



US011521294B2

(12) **United States Patent**
Ray et al.

(10) **Patent No.:** **US 11,521,294 B2**
(45) **Date of Patent:** **Dec. 6, 2022**

(54) **EFFICIENT MERGING OF ATOMIC OPERATIONS AT COMPUTING DEVICES**

(56) **References Cited**

(71) Applicant: **Intel Corporation**, Santa Clara, CA (US)

U.S. PATENT DOCUMENTS

(72) Inventors: **Joydeep Ray**, Folsom, CA (US); **Altug Koker**, El Dorado Hills, CA (US); **Abhishek R. Appu**, El Dorado Hills, CA (US); **Balaji Vembu**, Folsom, CA (US)

7,873,812 B1 1/2011 Mimar
8,108,610 B1* 1/2012 Glasco G06F 15/8015
711/118

10,325,344 B2 6/2019 Ray et al.
10,528,864 B2 1/2020 Dally et al.
10,796,401 B2 10/2020 Ray et al.
10,860,922 B2 12/2020 Dally et al.

(Continued)

(73) Assignee: **INTEL CORPORATION**, Santa Clara, CA (US)

OTHER PUBLICATIONS

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

Goodfellow, et al. "Adaptive Computation and Machine Learning Series", Book, Nov. 18, 2016, pp. 98-165, Chapter 5, The MIT Press, Cambridge, MA.

(Continued)

(21) Appl. No.: **17/036,950**

Primary Examiner — Jennifer Mehmood

(22) Filed: **Sep. 29, 2020**

Assistant Examiner — Donna J. Ricks

(65) **Prior Publication Data**

US 2021/0065329 A1 Mar. 4, 2021

(74) *Attorney, Agent, or Firm* — Jaffery Watson

Mendonsa & Hamilton LLP

Related U.S. Application Data

(57) **ABSTRACT**

(63) Continuation of application No. 16/414,251, filed on May 16, 2019, now Pat. No. 10,796,401, which is a continuation of application No. 15/488,565, filed on Apr. 17, 2017, now Pat. No. 10,325,344.

A mechanism is described for facilitating dynamic merging of atomic operations in computing devices. A method of embodiments, as described herein, includes facilitating detecting atomic messages and a plurality of slot addresses. The method further includes comparing one or more slot addresses of the plurality of slot addresses with other slot addresses of the plurality of slot addresses to seek one or more matched slot addresses, where the one or more matched slot addresses are merged into one or more merged groups. The method may further include generating one or more merged atomic operations based on and corresponding to the one or more merged groups.

(51) **Int. Cl.**

G06T 1/60 (2006.01)

G06T 1/20 (2006.01)

(52) **U.S. Cl.**

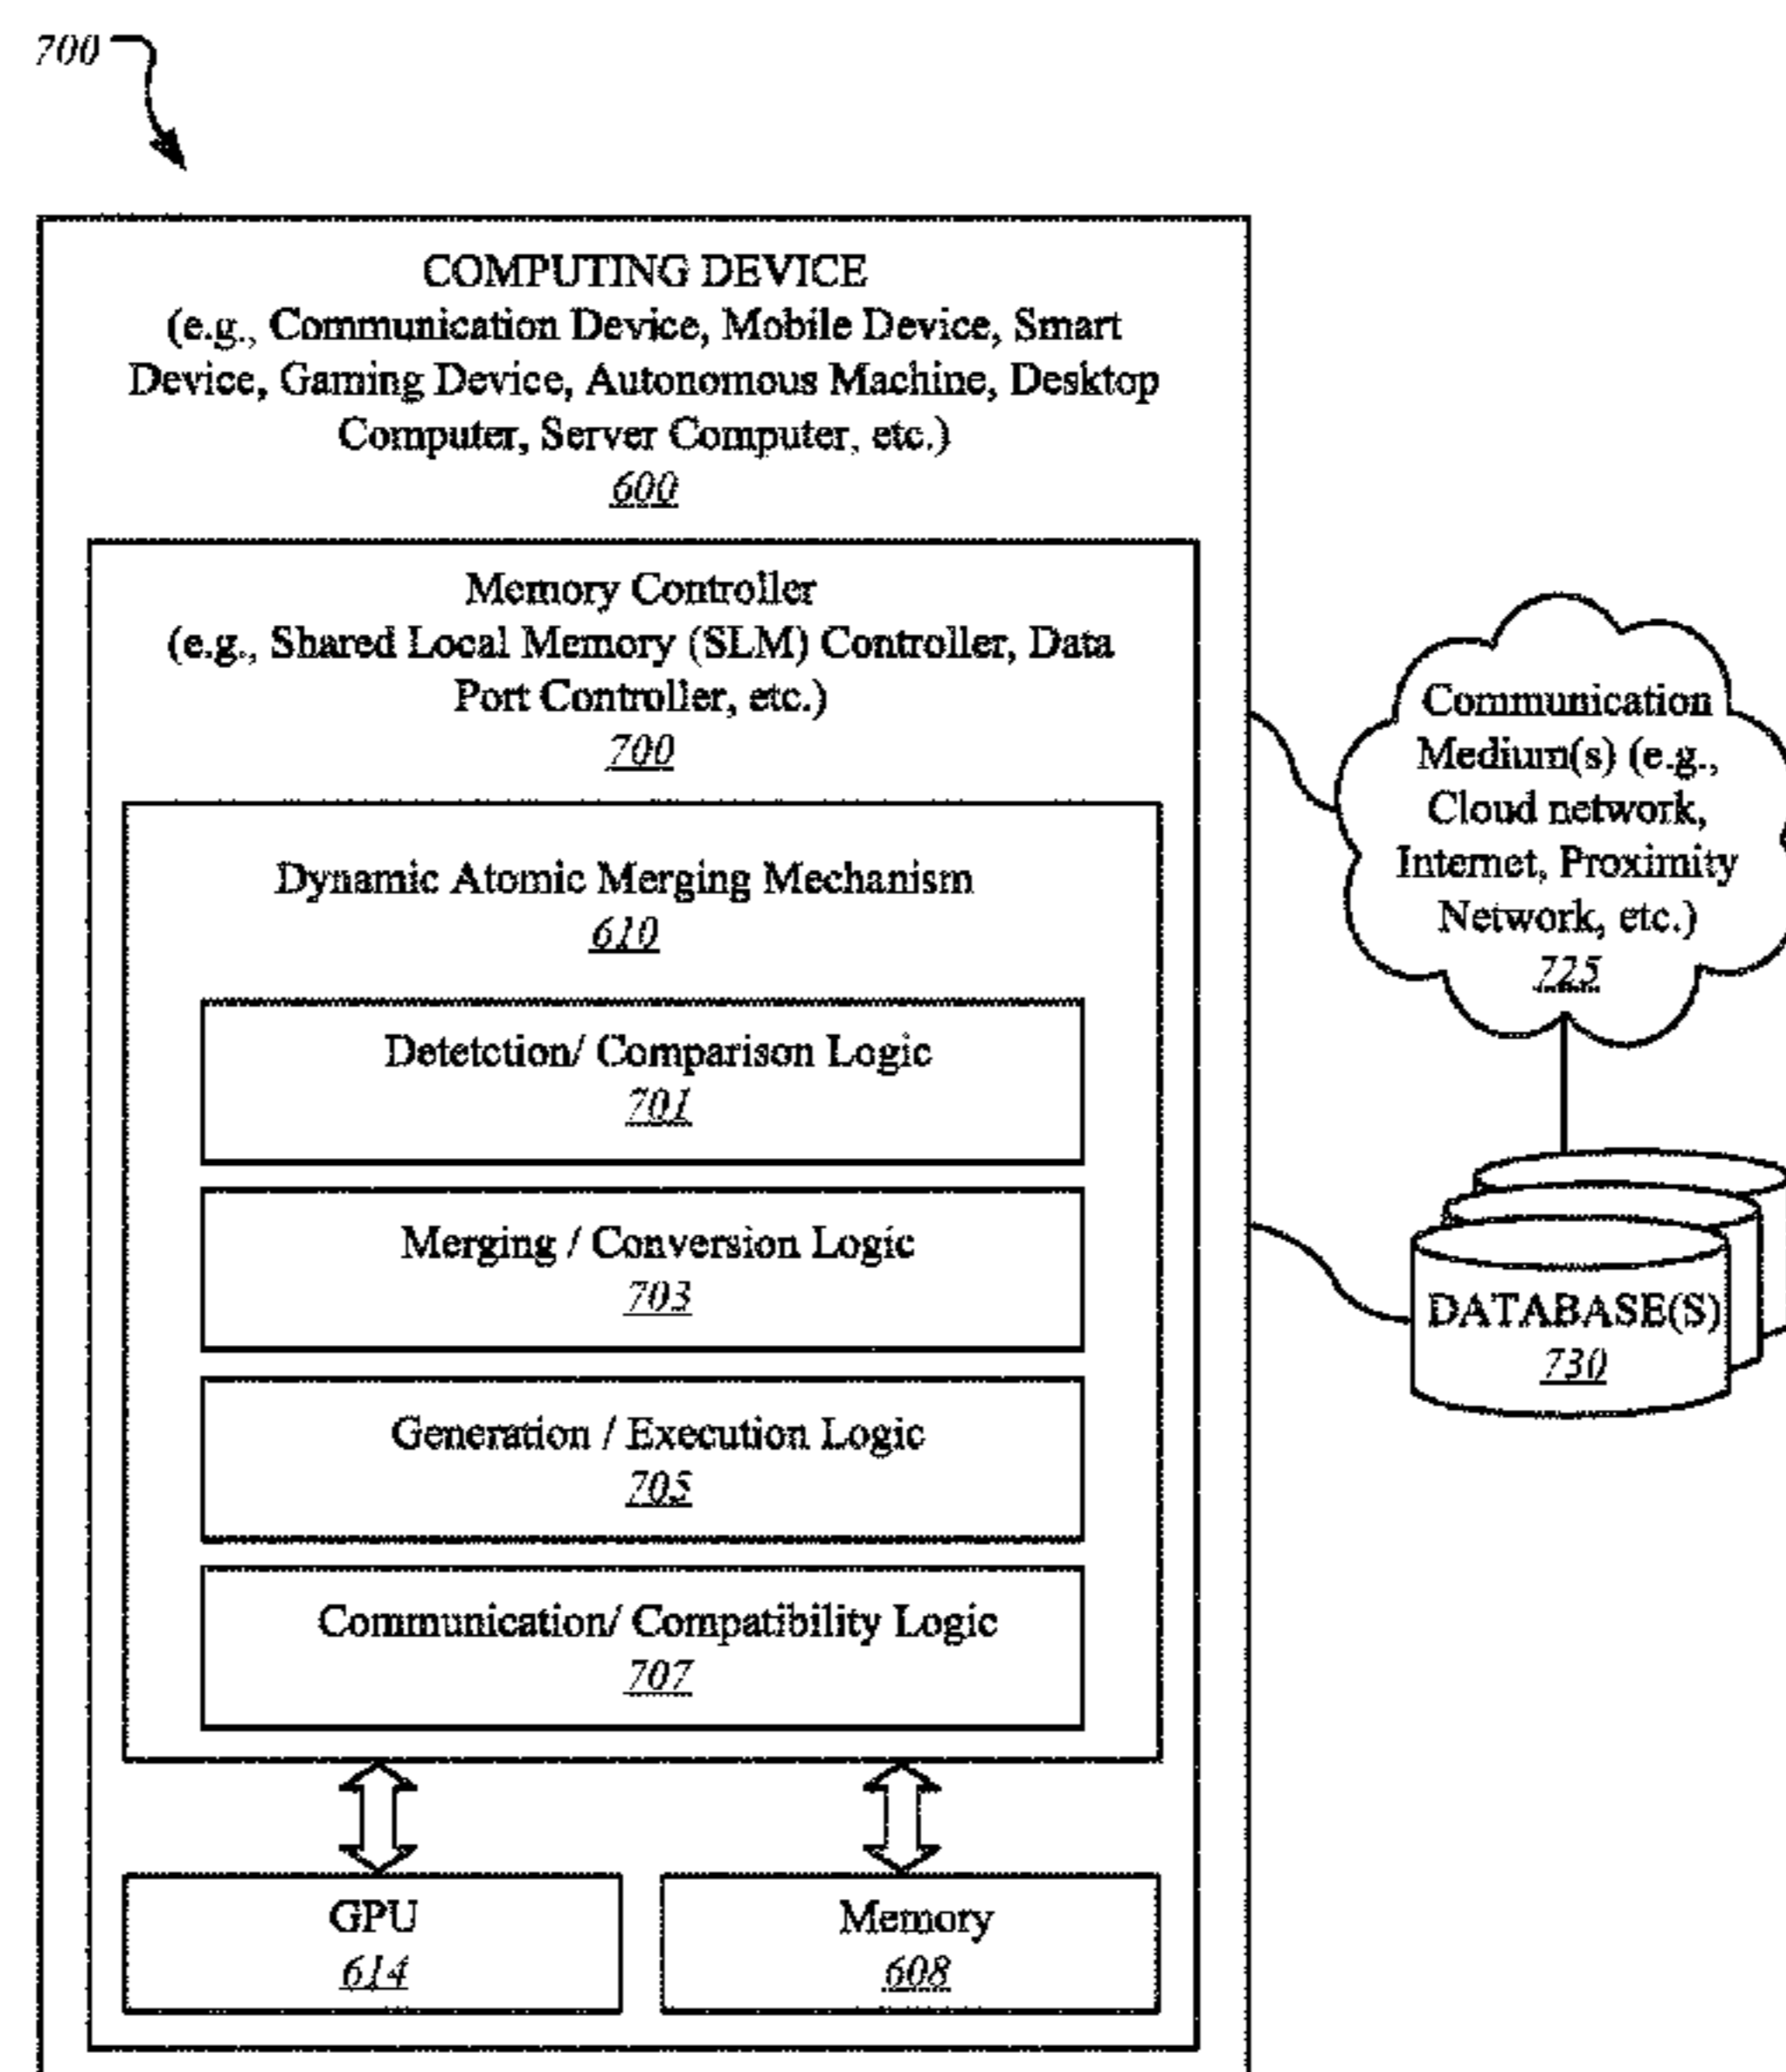
CPC . **G06T 1/60** (2013.01); **G06T 1/20** (2013.01)

(58) **Field of Classification Search**

CPC G06T 1/60; G06T 1/20

See application file for complete search history.

25 Claims, 34 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

10,891,538 B2 1/2021 Dally et al.
2005/0235134 A1 10/2005 O'Sullivan
2012/0023316 A1 1/2012 Flachs et al.
2013/0232322 A1* 9/2013 Fetterman G06F 9/3887
712/225
2014/0181421 A1 6/2014 O'Connor et al.
2014/0366033 A1* 12/2014 Nystad G06F 9/50
718/104
2016/0062947 A1 3/2016 Chetlur et al.
2016/0350262 A1 12/2016 Sarangi et al.
2018/0046906 A1 2/2018 Dally et al.
2018/0300846 A1 10/2018 Ray et al.
2019/0347759 A1 11/2019 Ray et al.

OTHER PUBLICATIONS

Ross, et al. "Intel Processor Graphics: Architecture & Programming", Power Point Presentation, Aug. 2015, 78 pages, Intel Corporation, Santa Clara, CA.
Shane Cook, "CUDA Programming", Book, 2013, pp. 37-52, Chapter 3, Elsevier Inc., Amsterdam Netherlands.
Nicholas Wilt, "The CUDA Handbook; A Comprehensive Guide to GPU Programming", Book, Jun. 22, 2013, pp. 41-57, Addison-Wesley Professional, Boston, MA.
Stephen Junking, "The Compute Architecture of Intel Processor Graphics Gen9", paper, Aug. 14, 2015, 22 pages, Version 1.0, Intel Corporation, Santa Clara, CA.

* cited by examiner

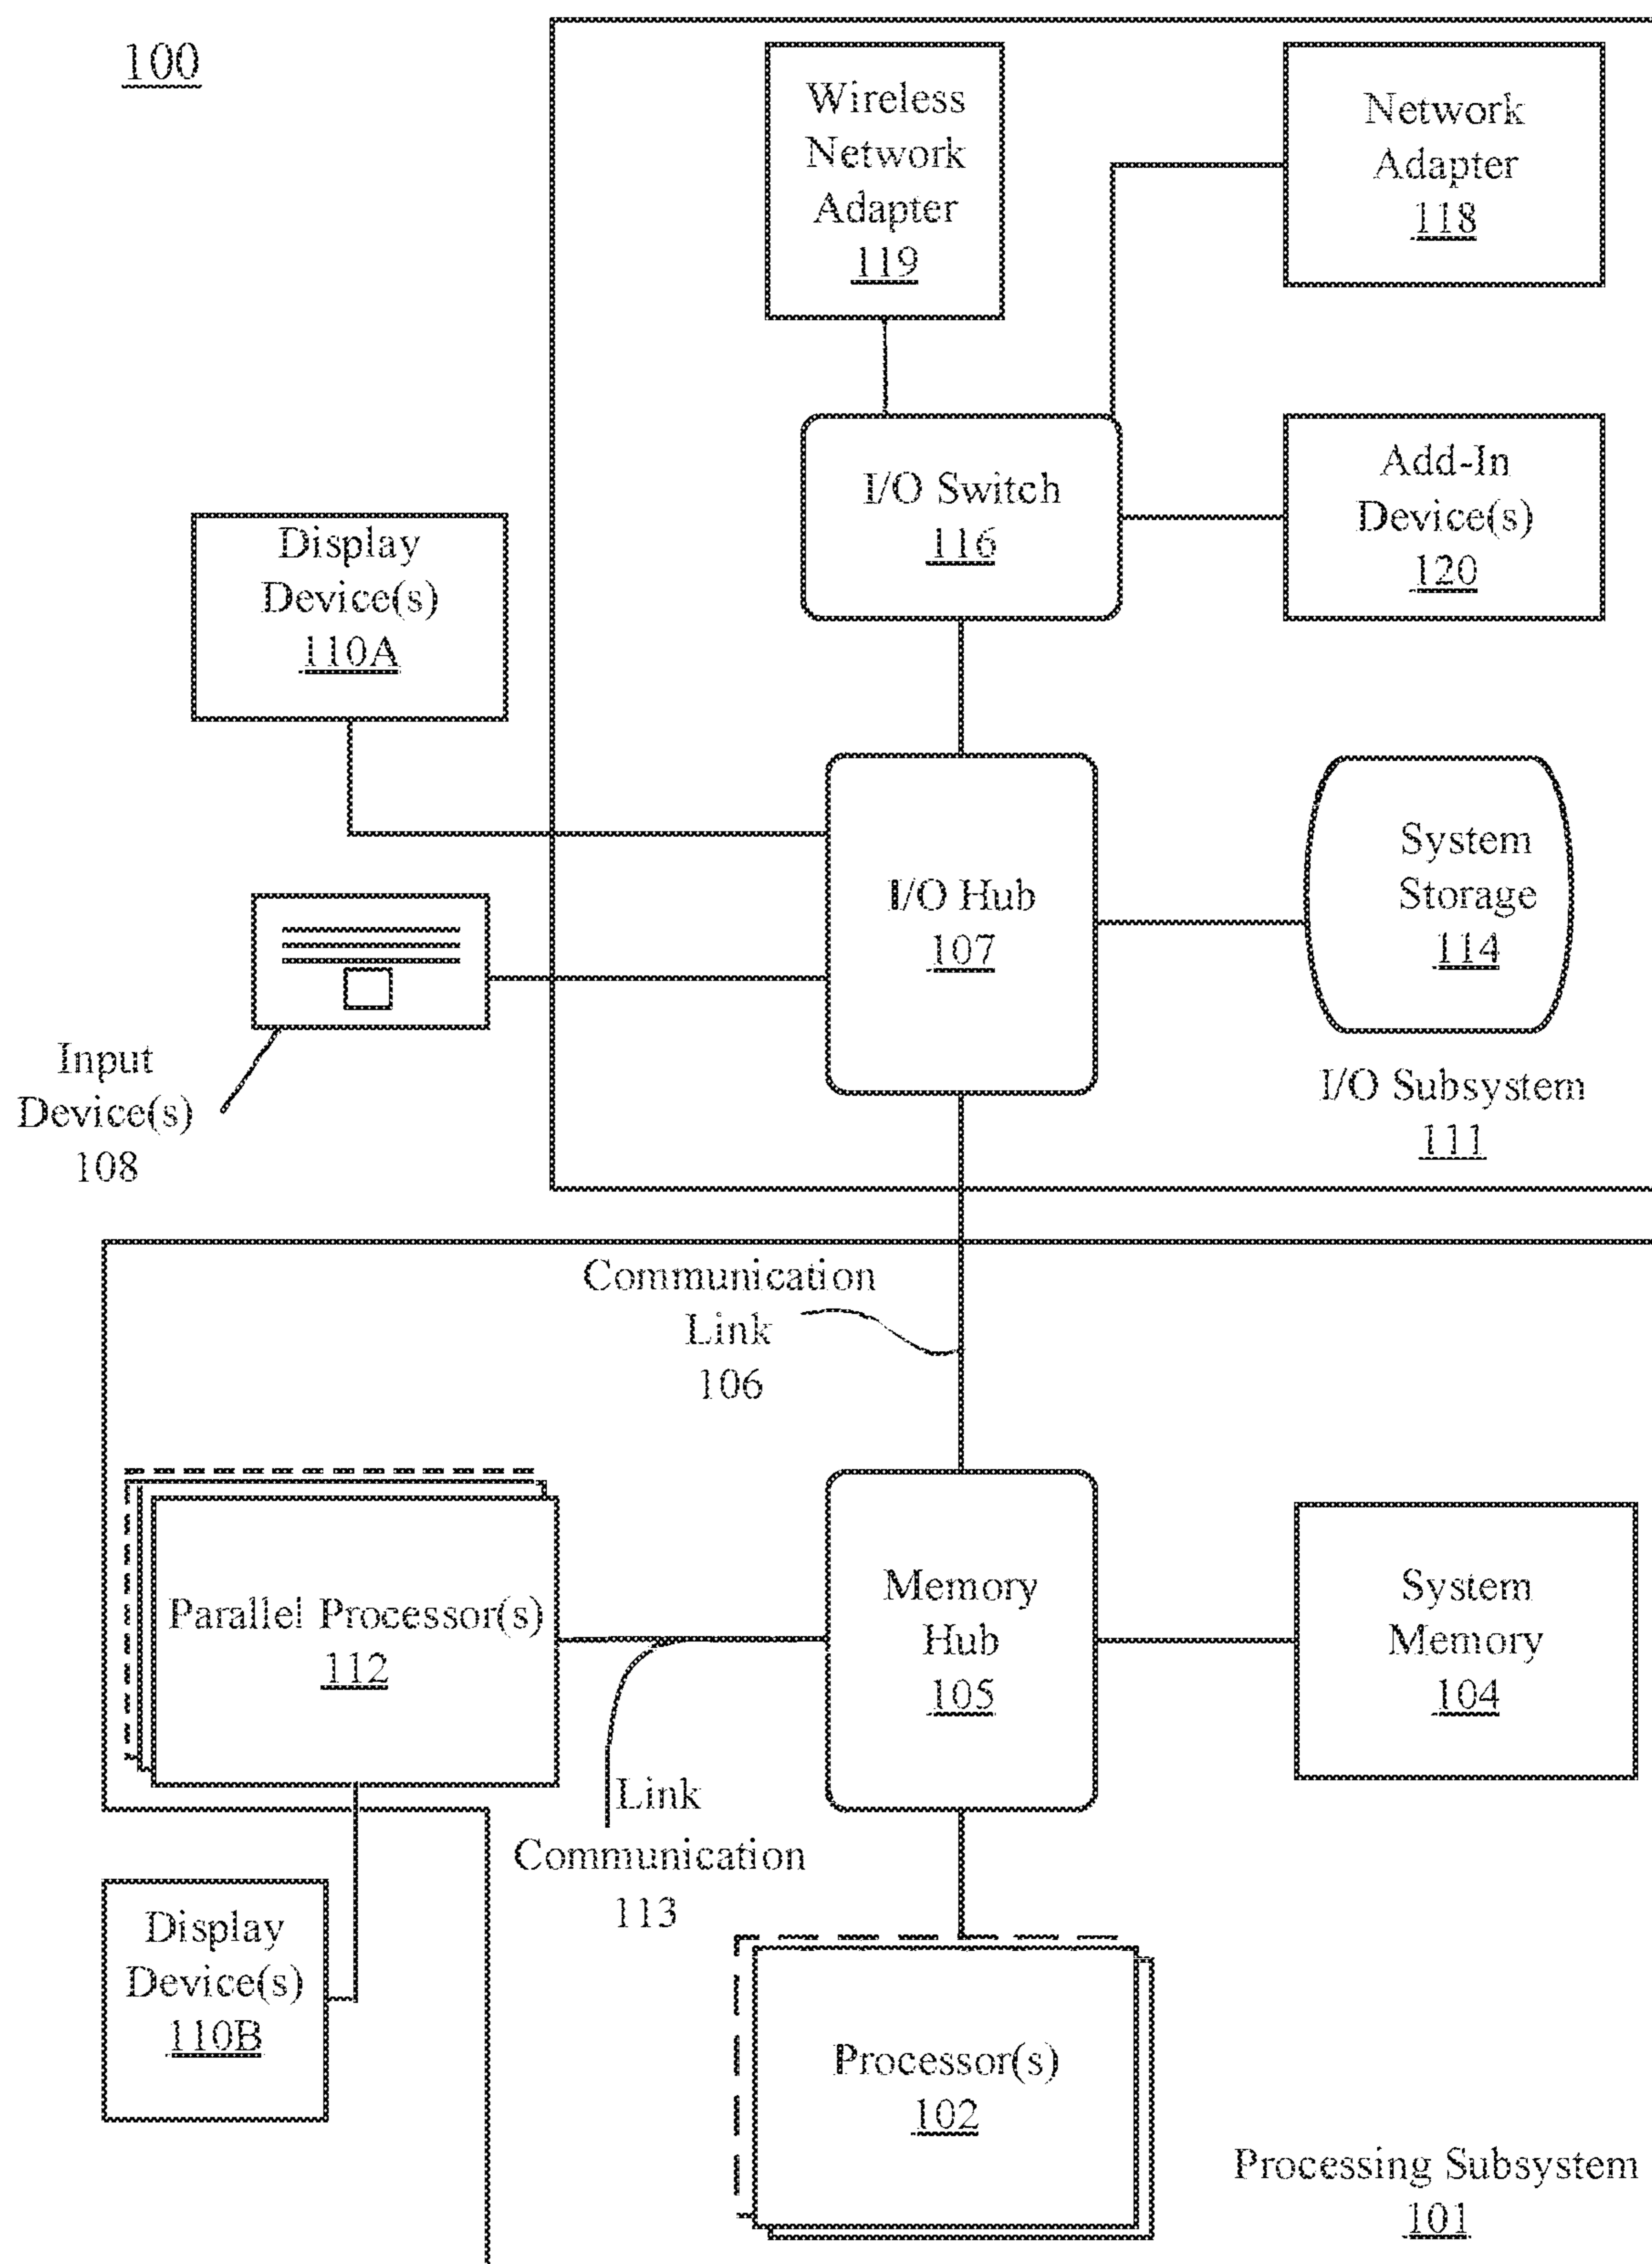


FIG. 1




FIG. 2A




FIG. 2B




FIG. 2C

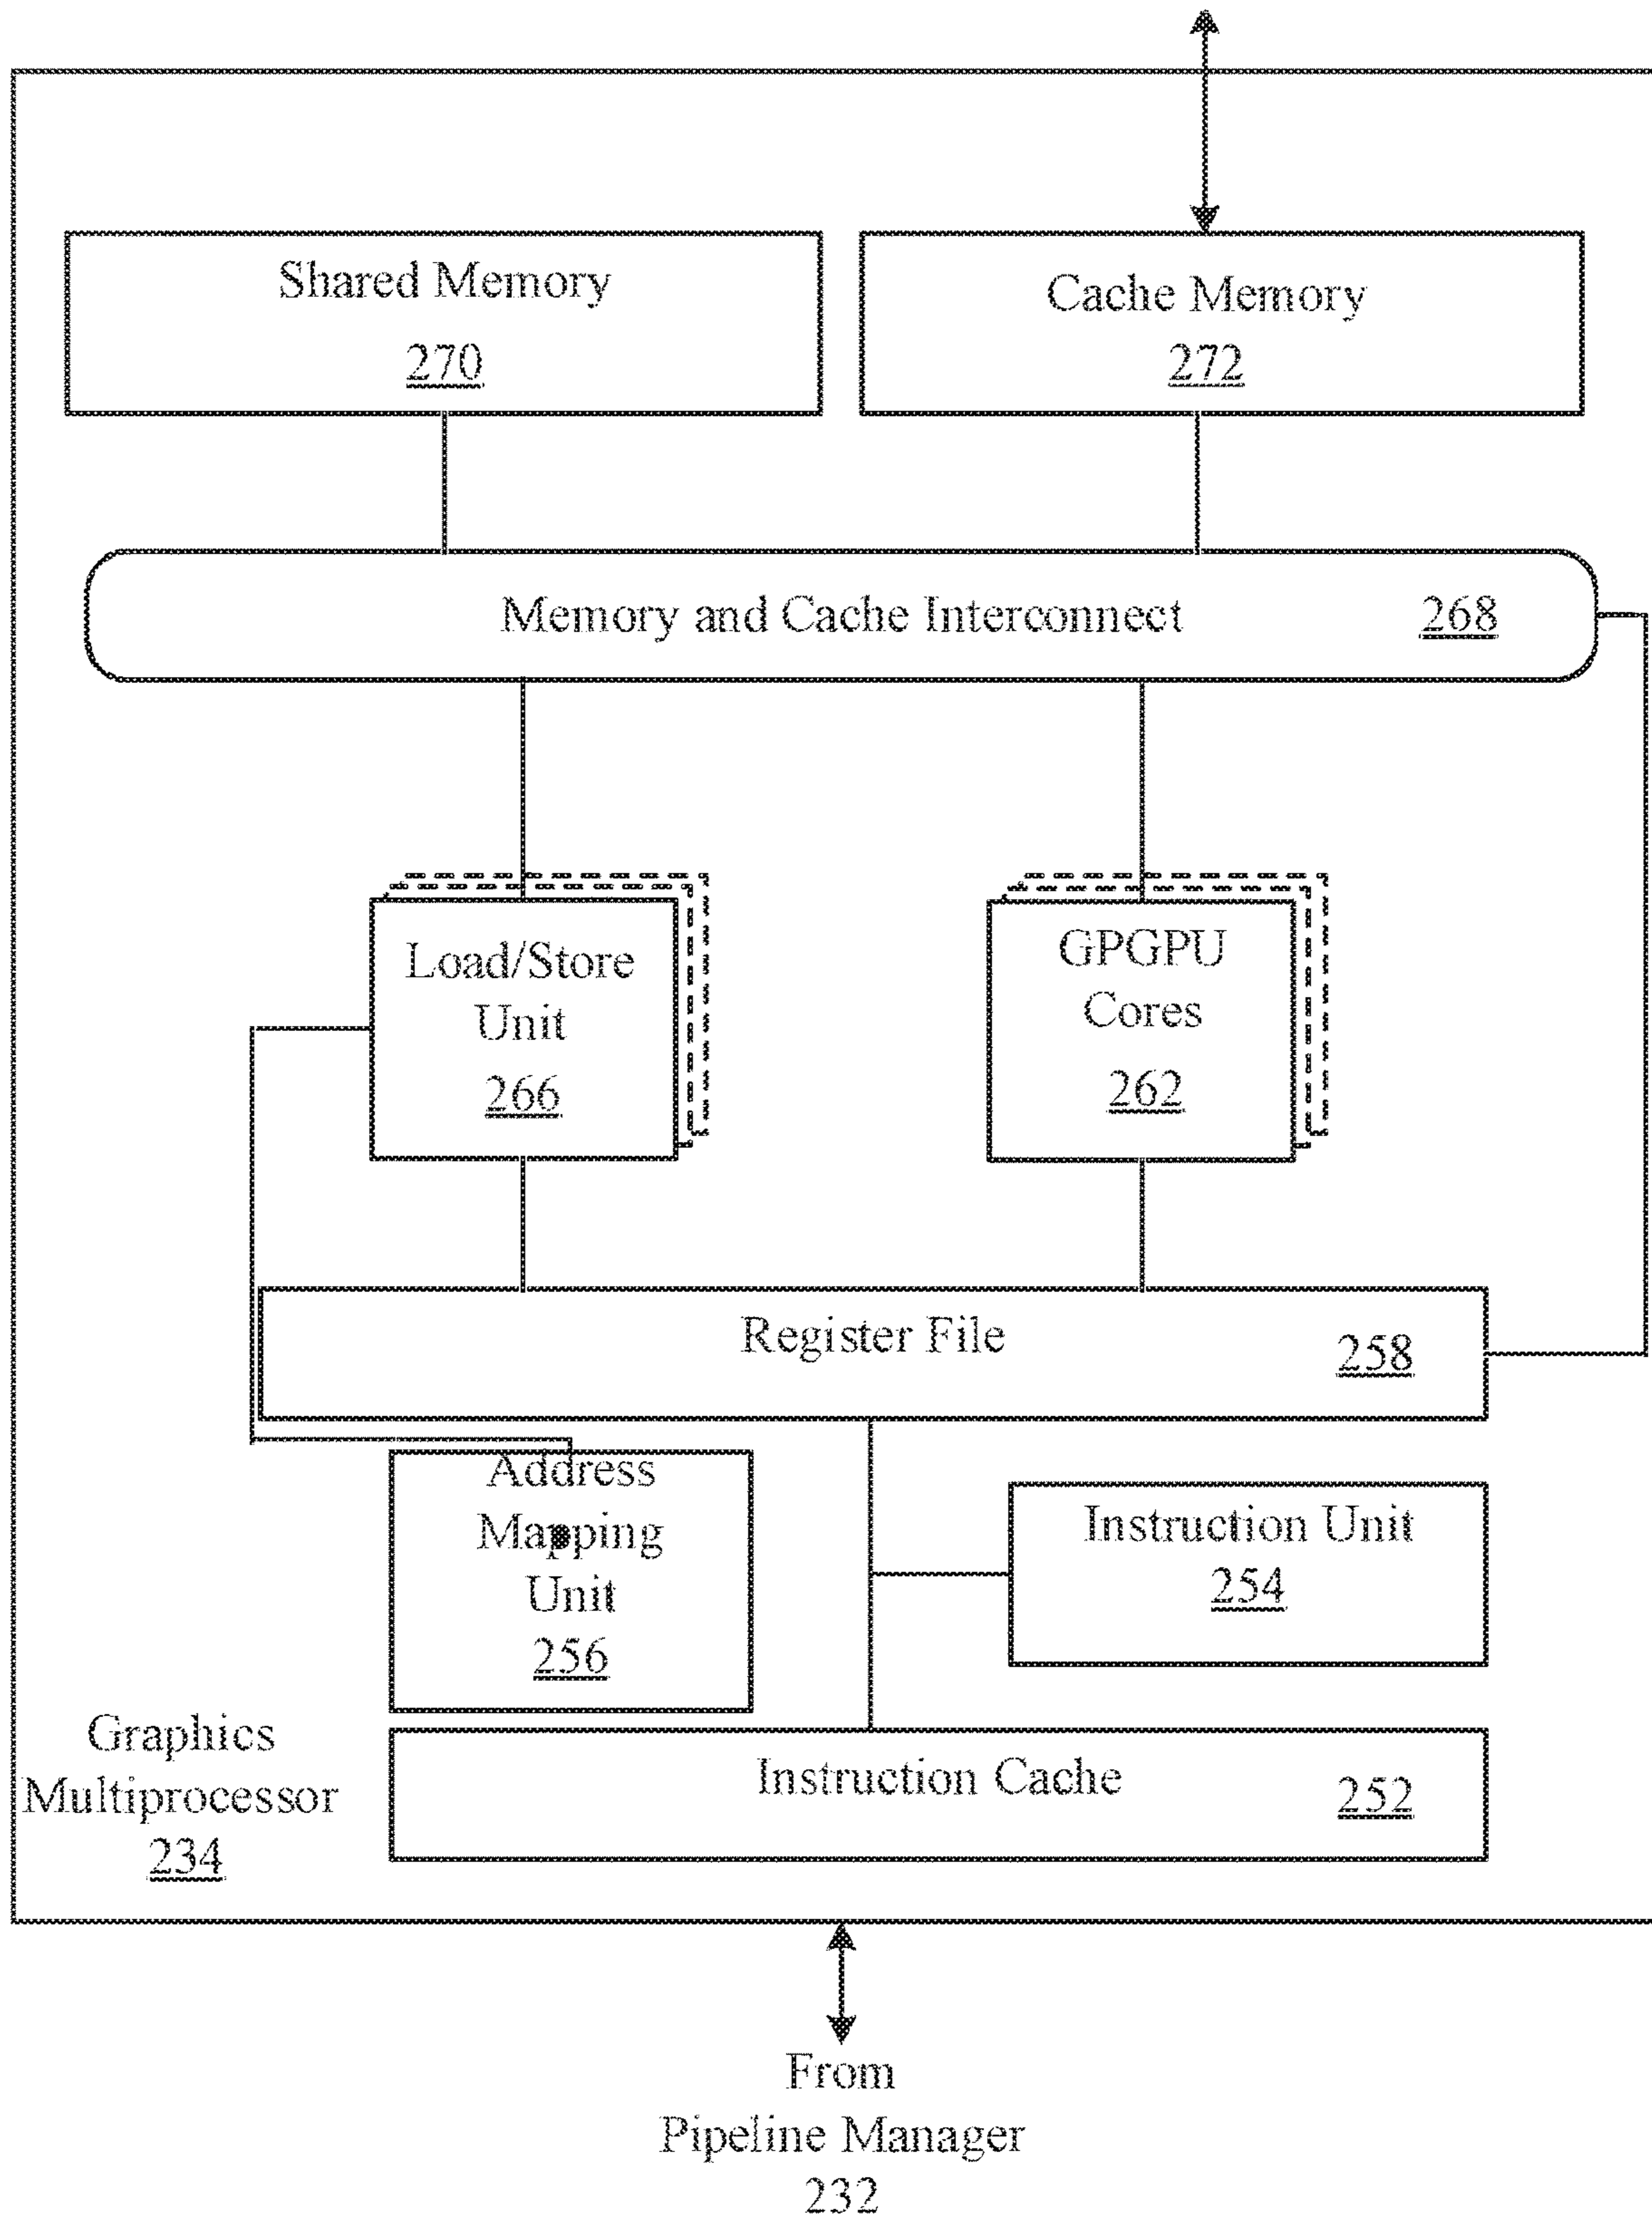


FIG. 2D

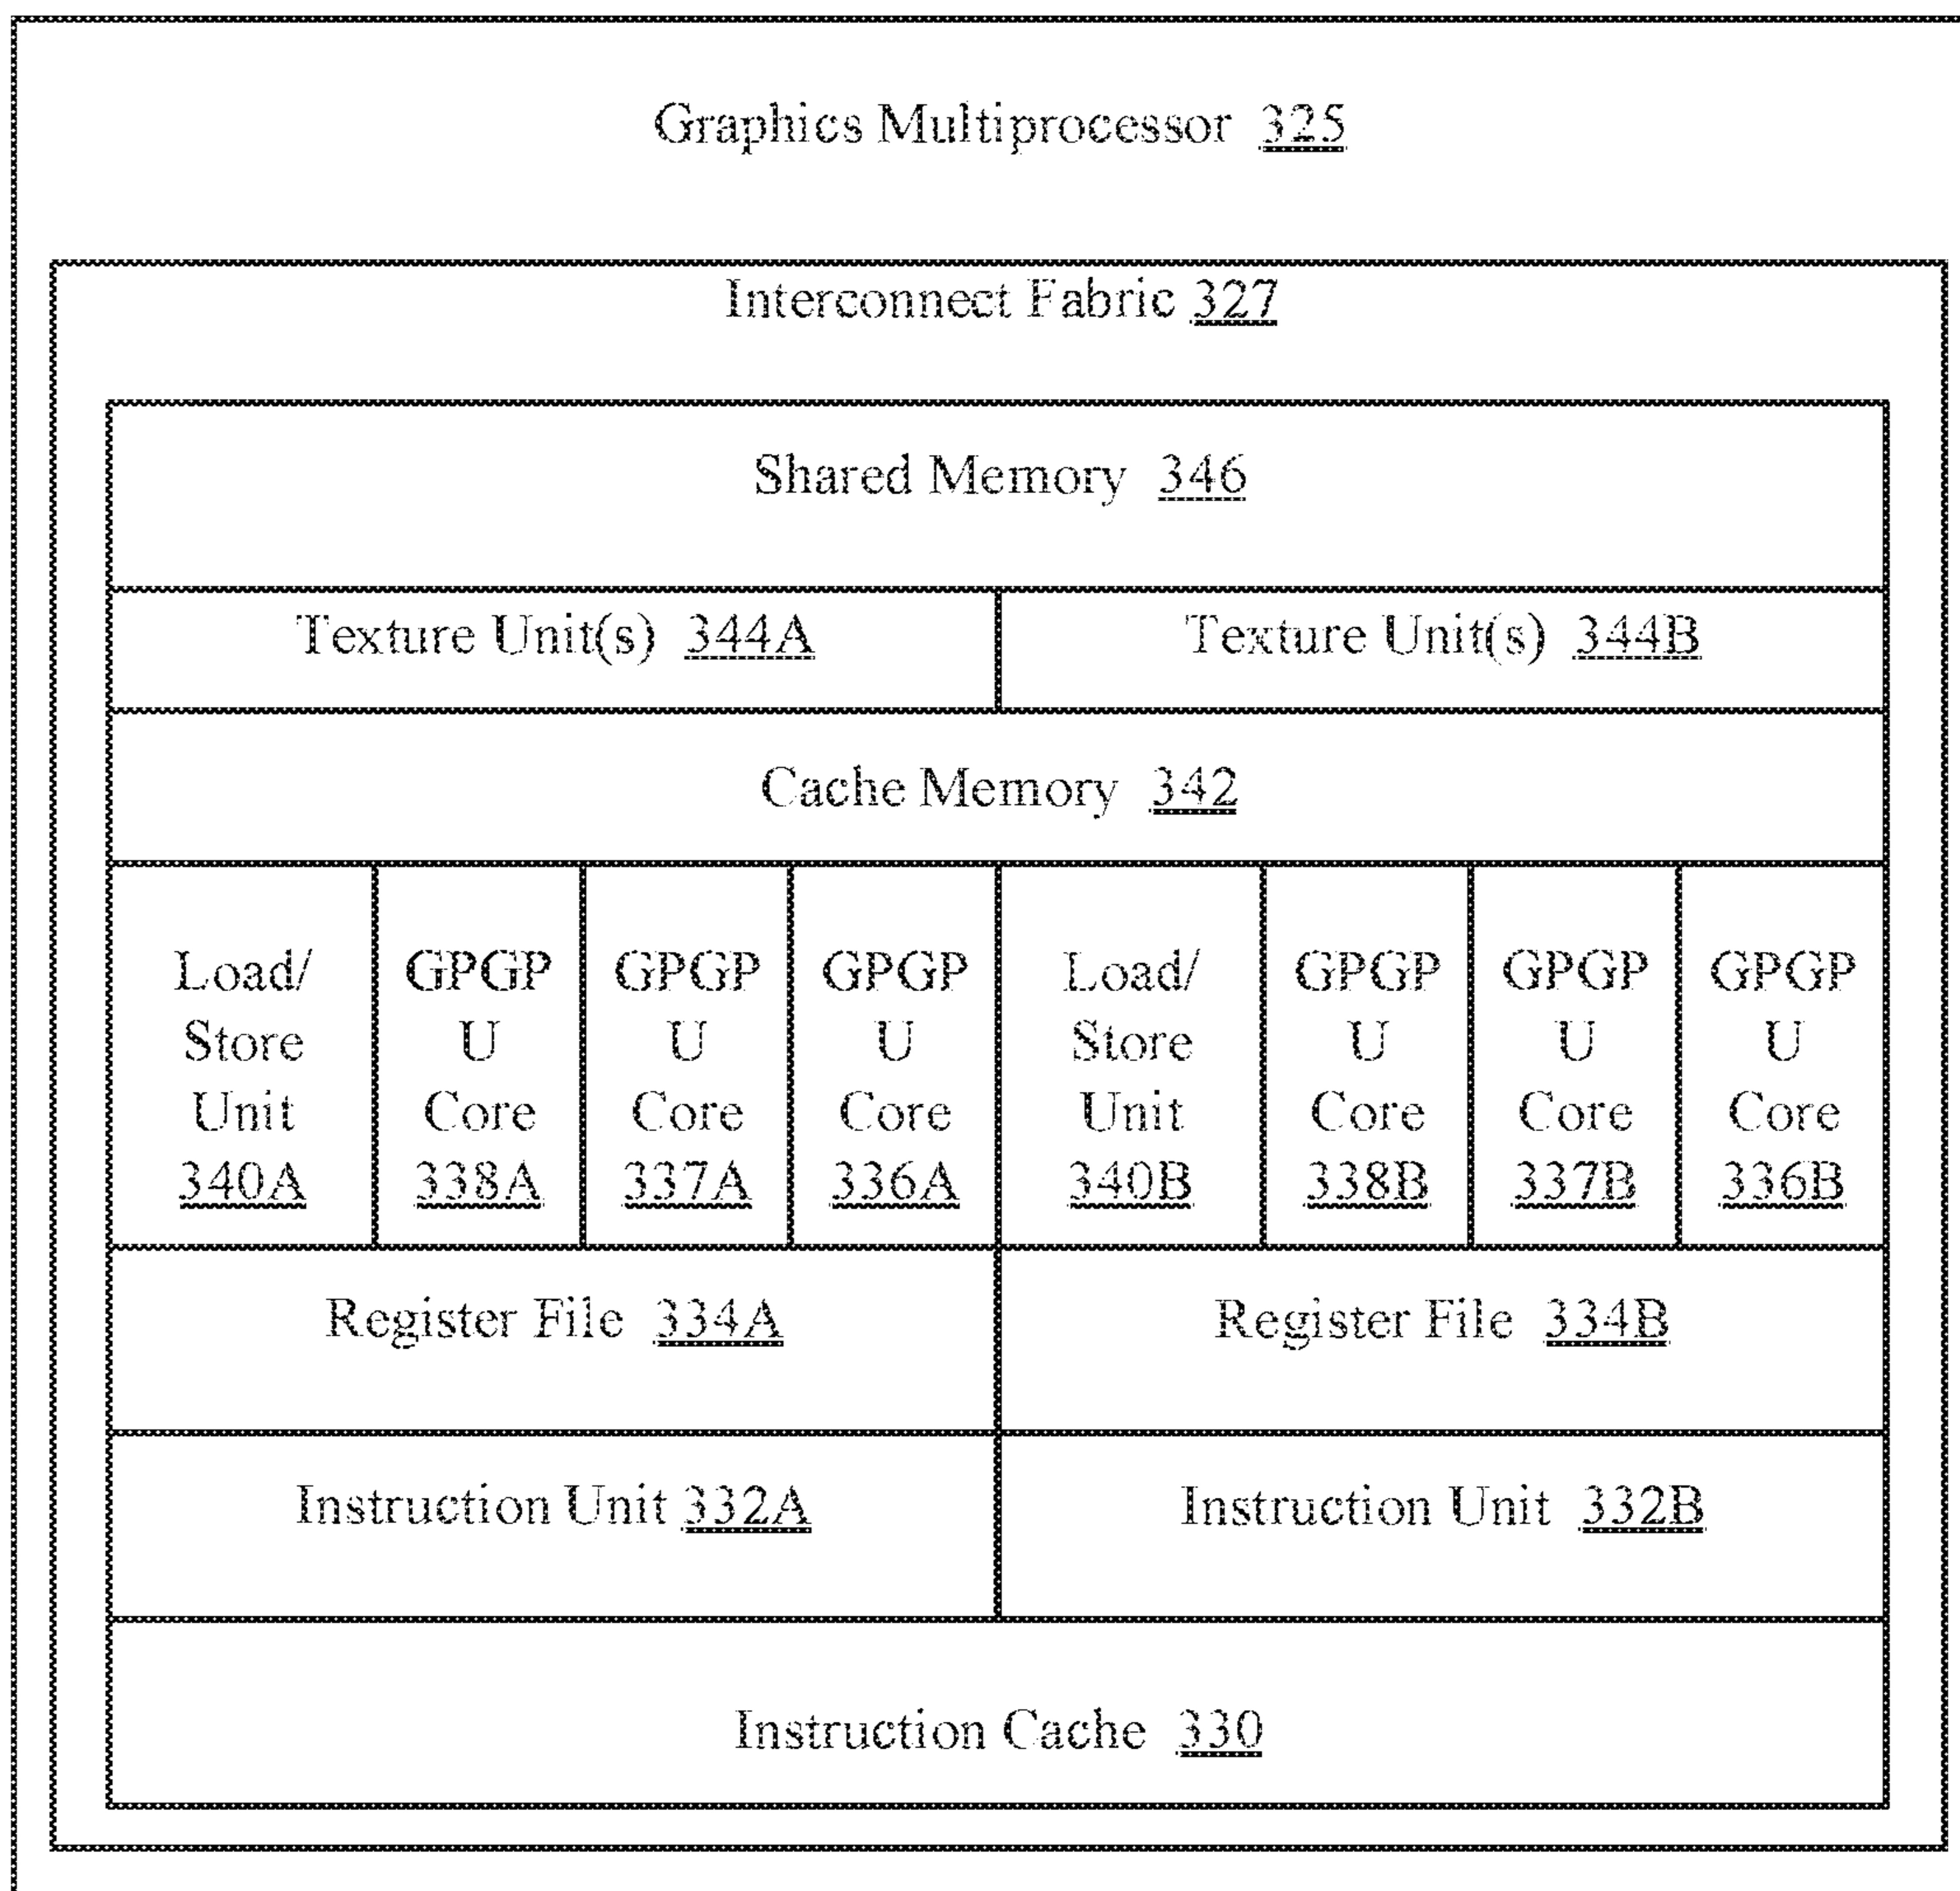


FIG. 3A

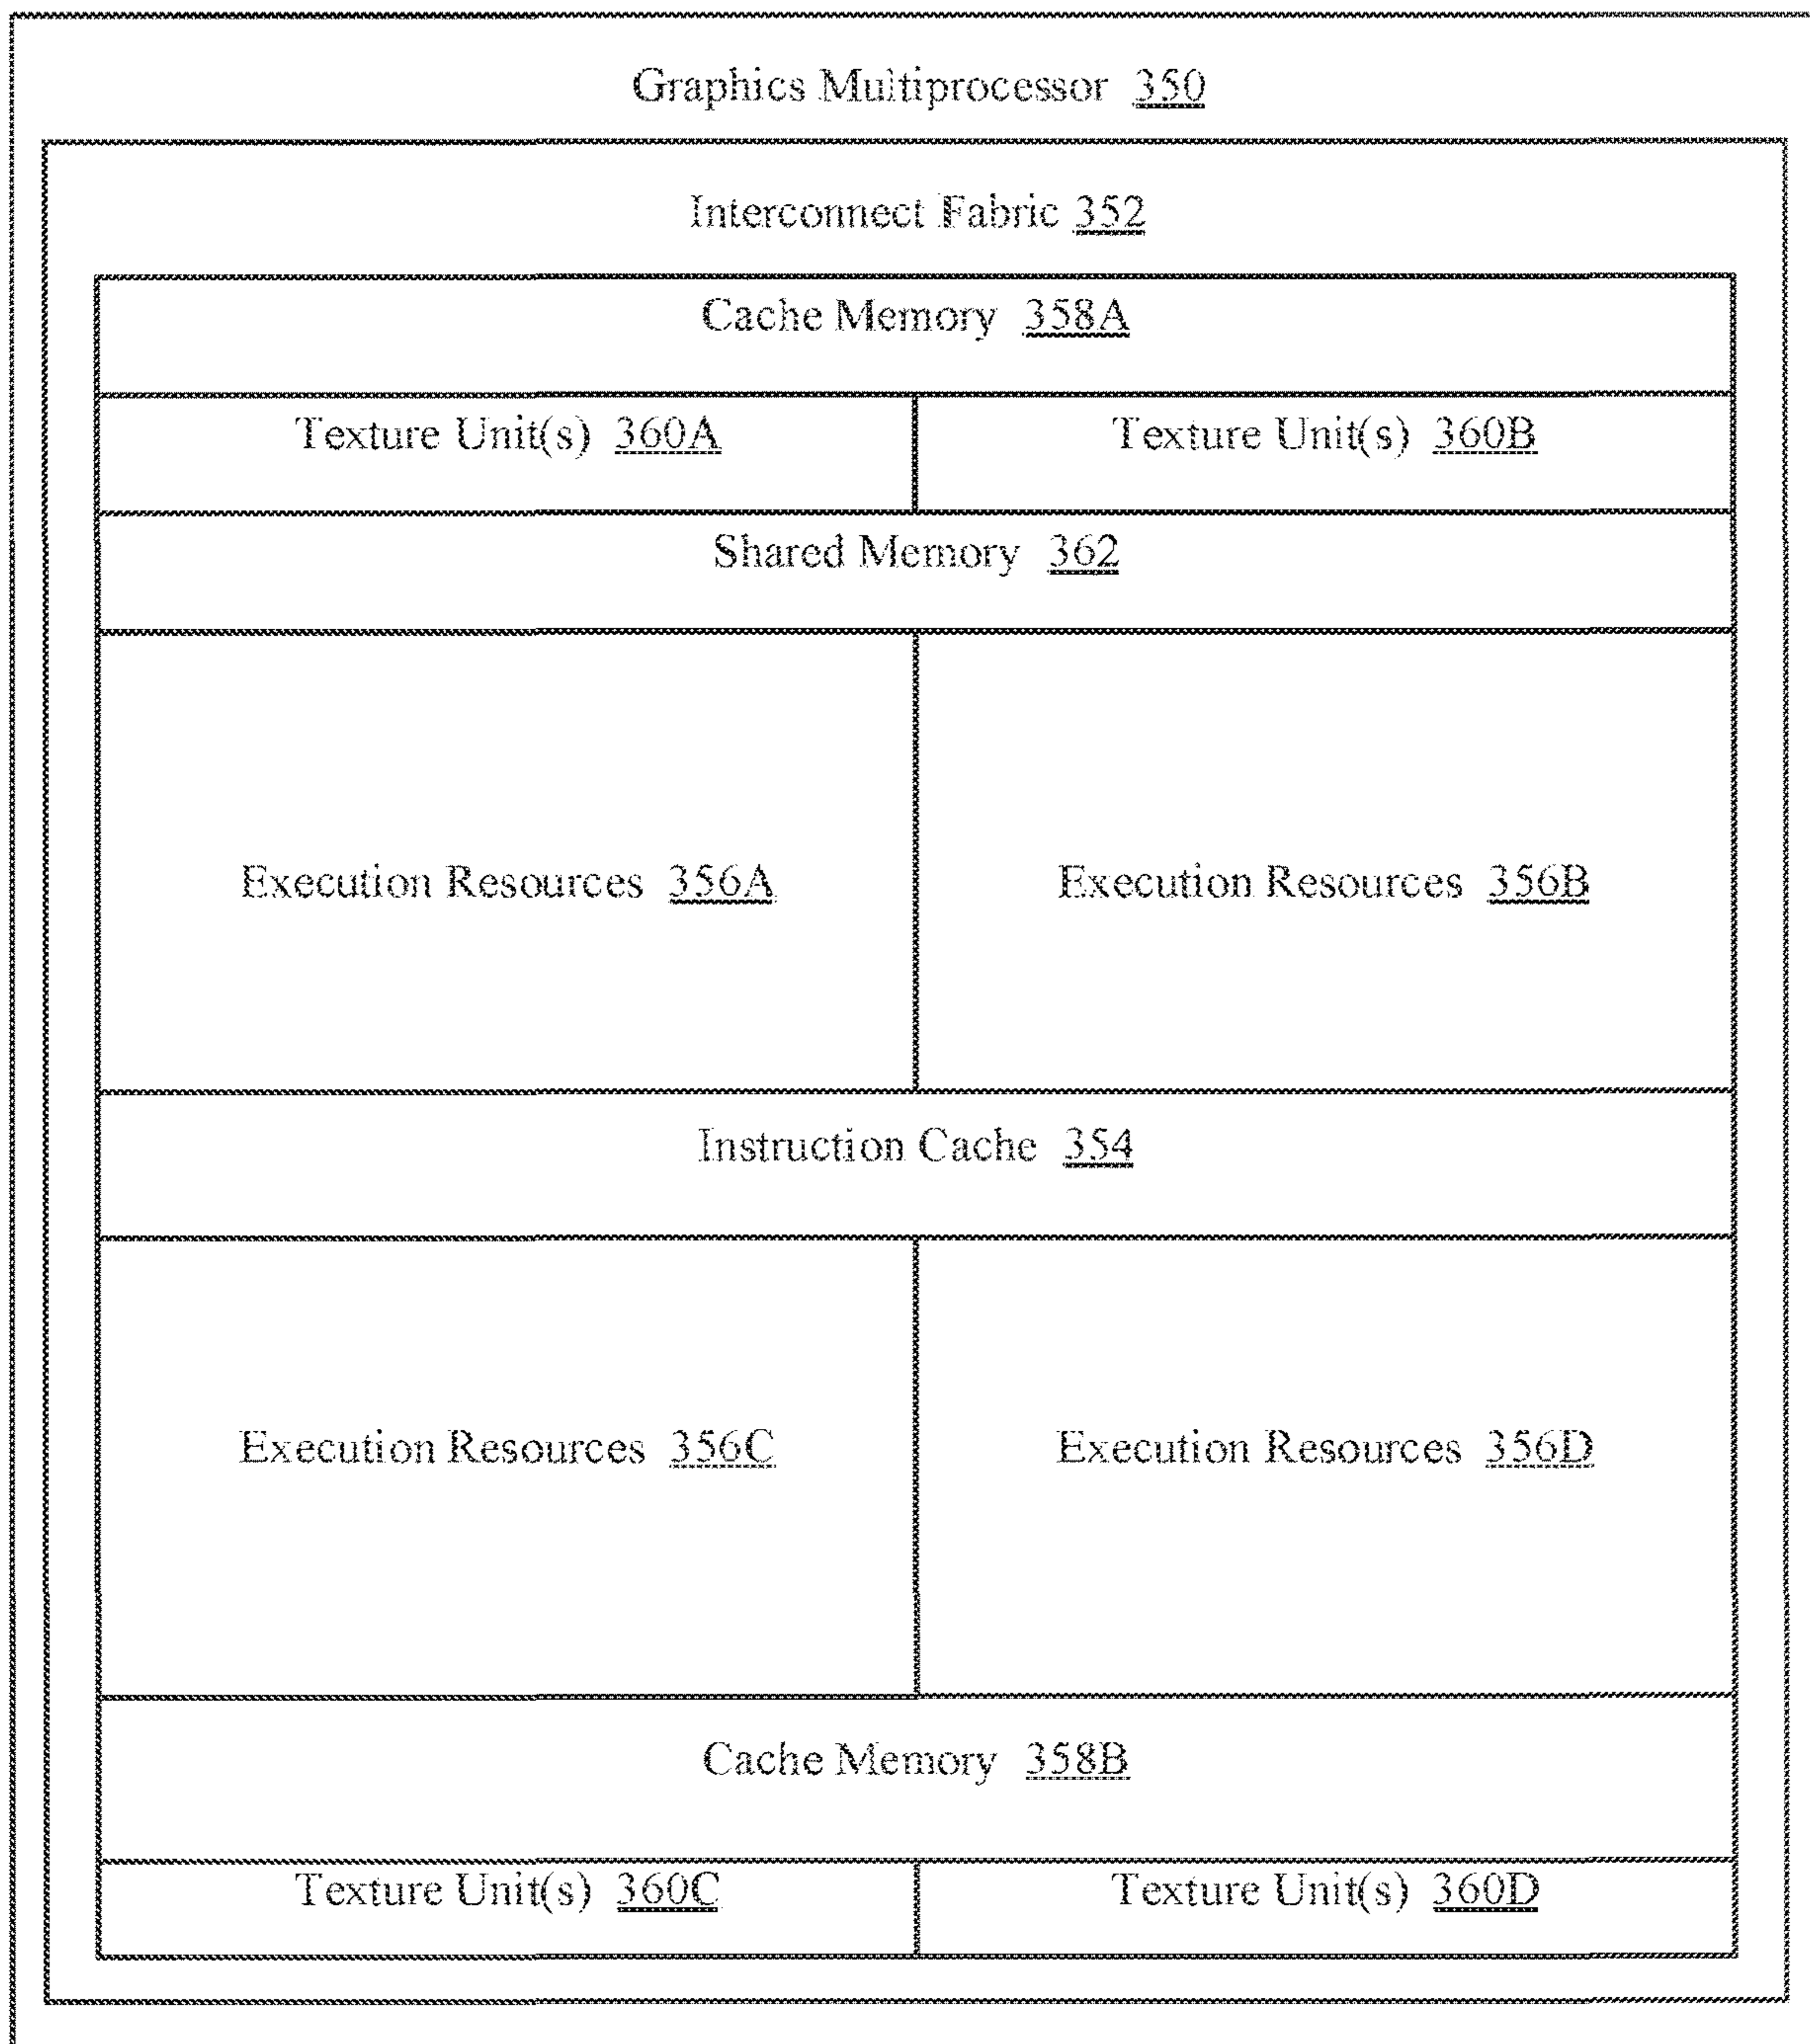


FIG. 3B


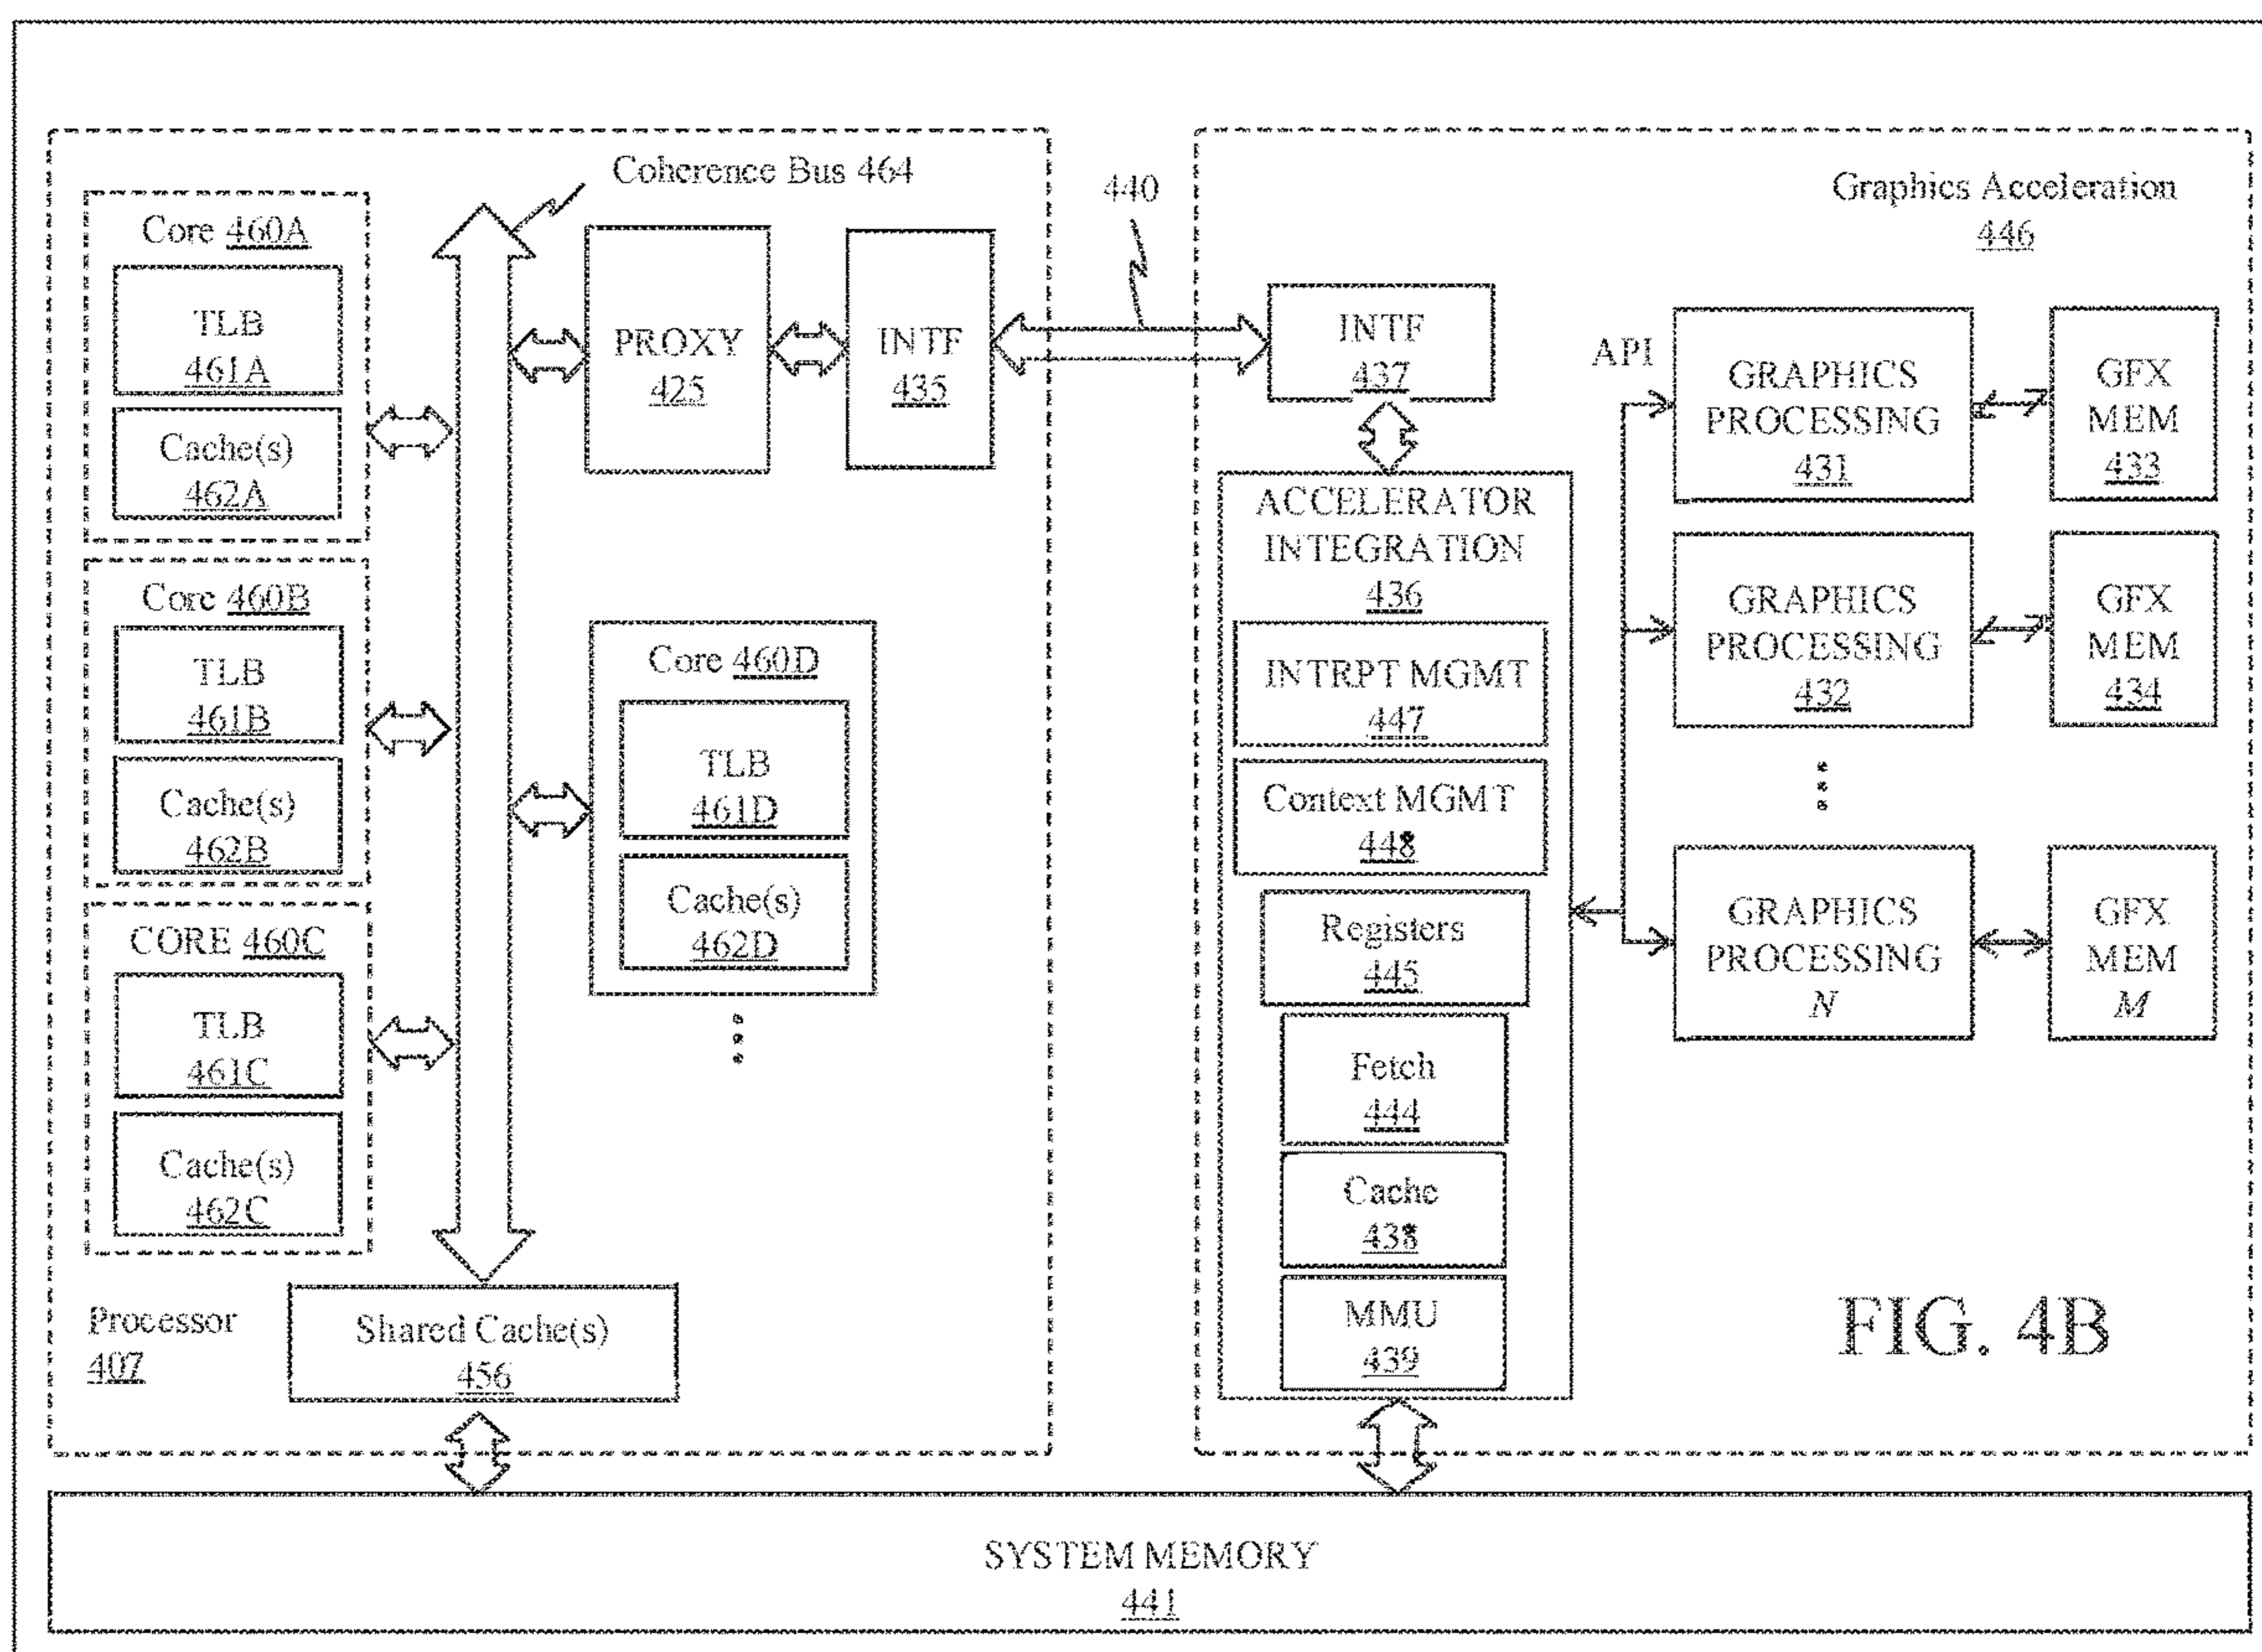
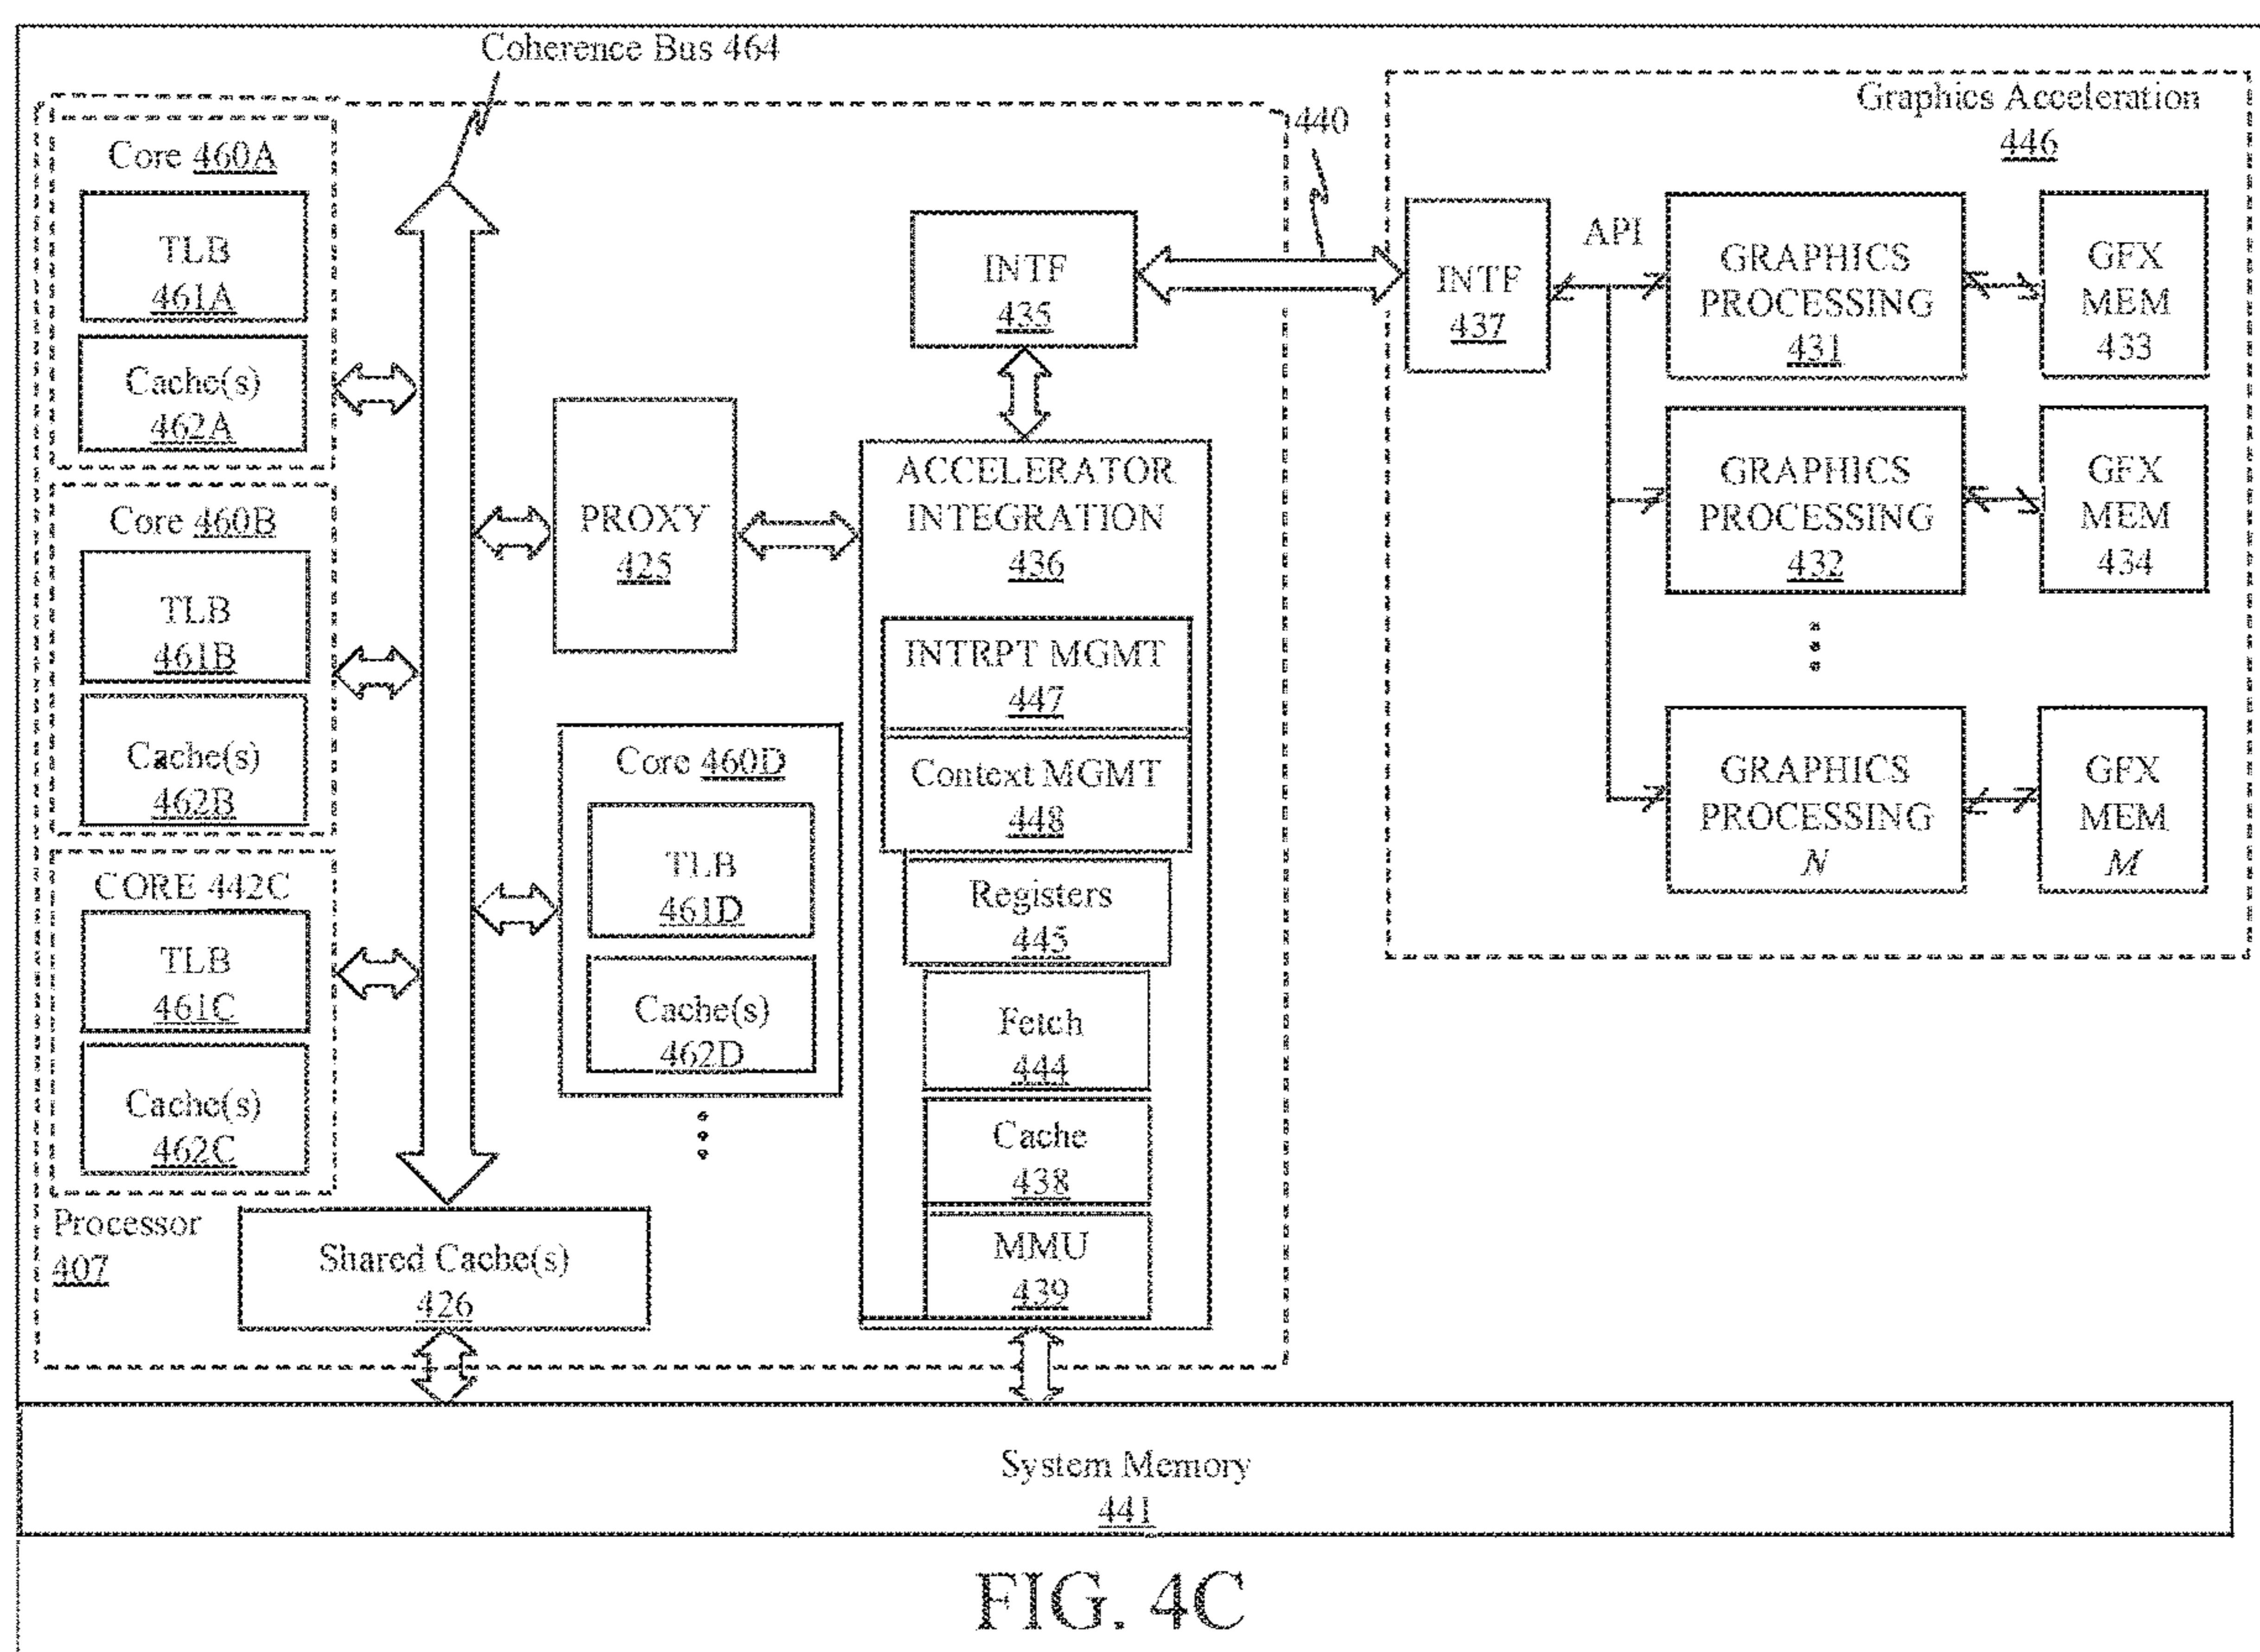


FIG. 4A





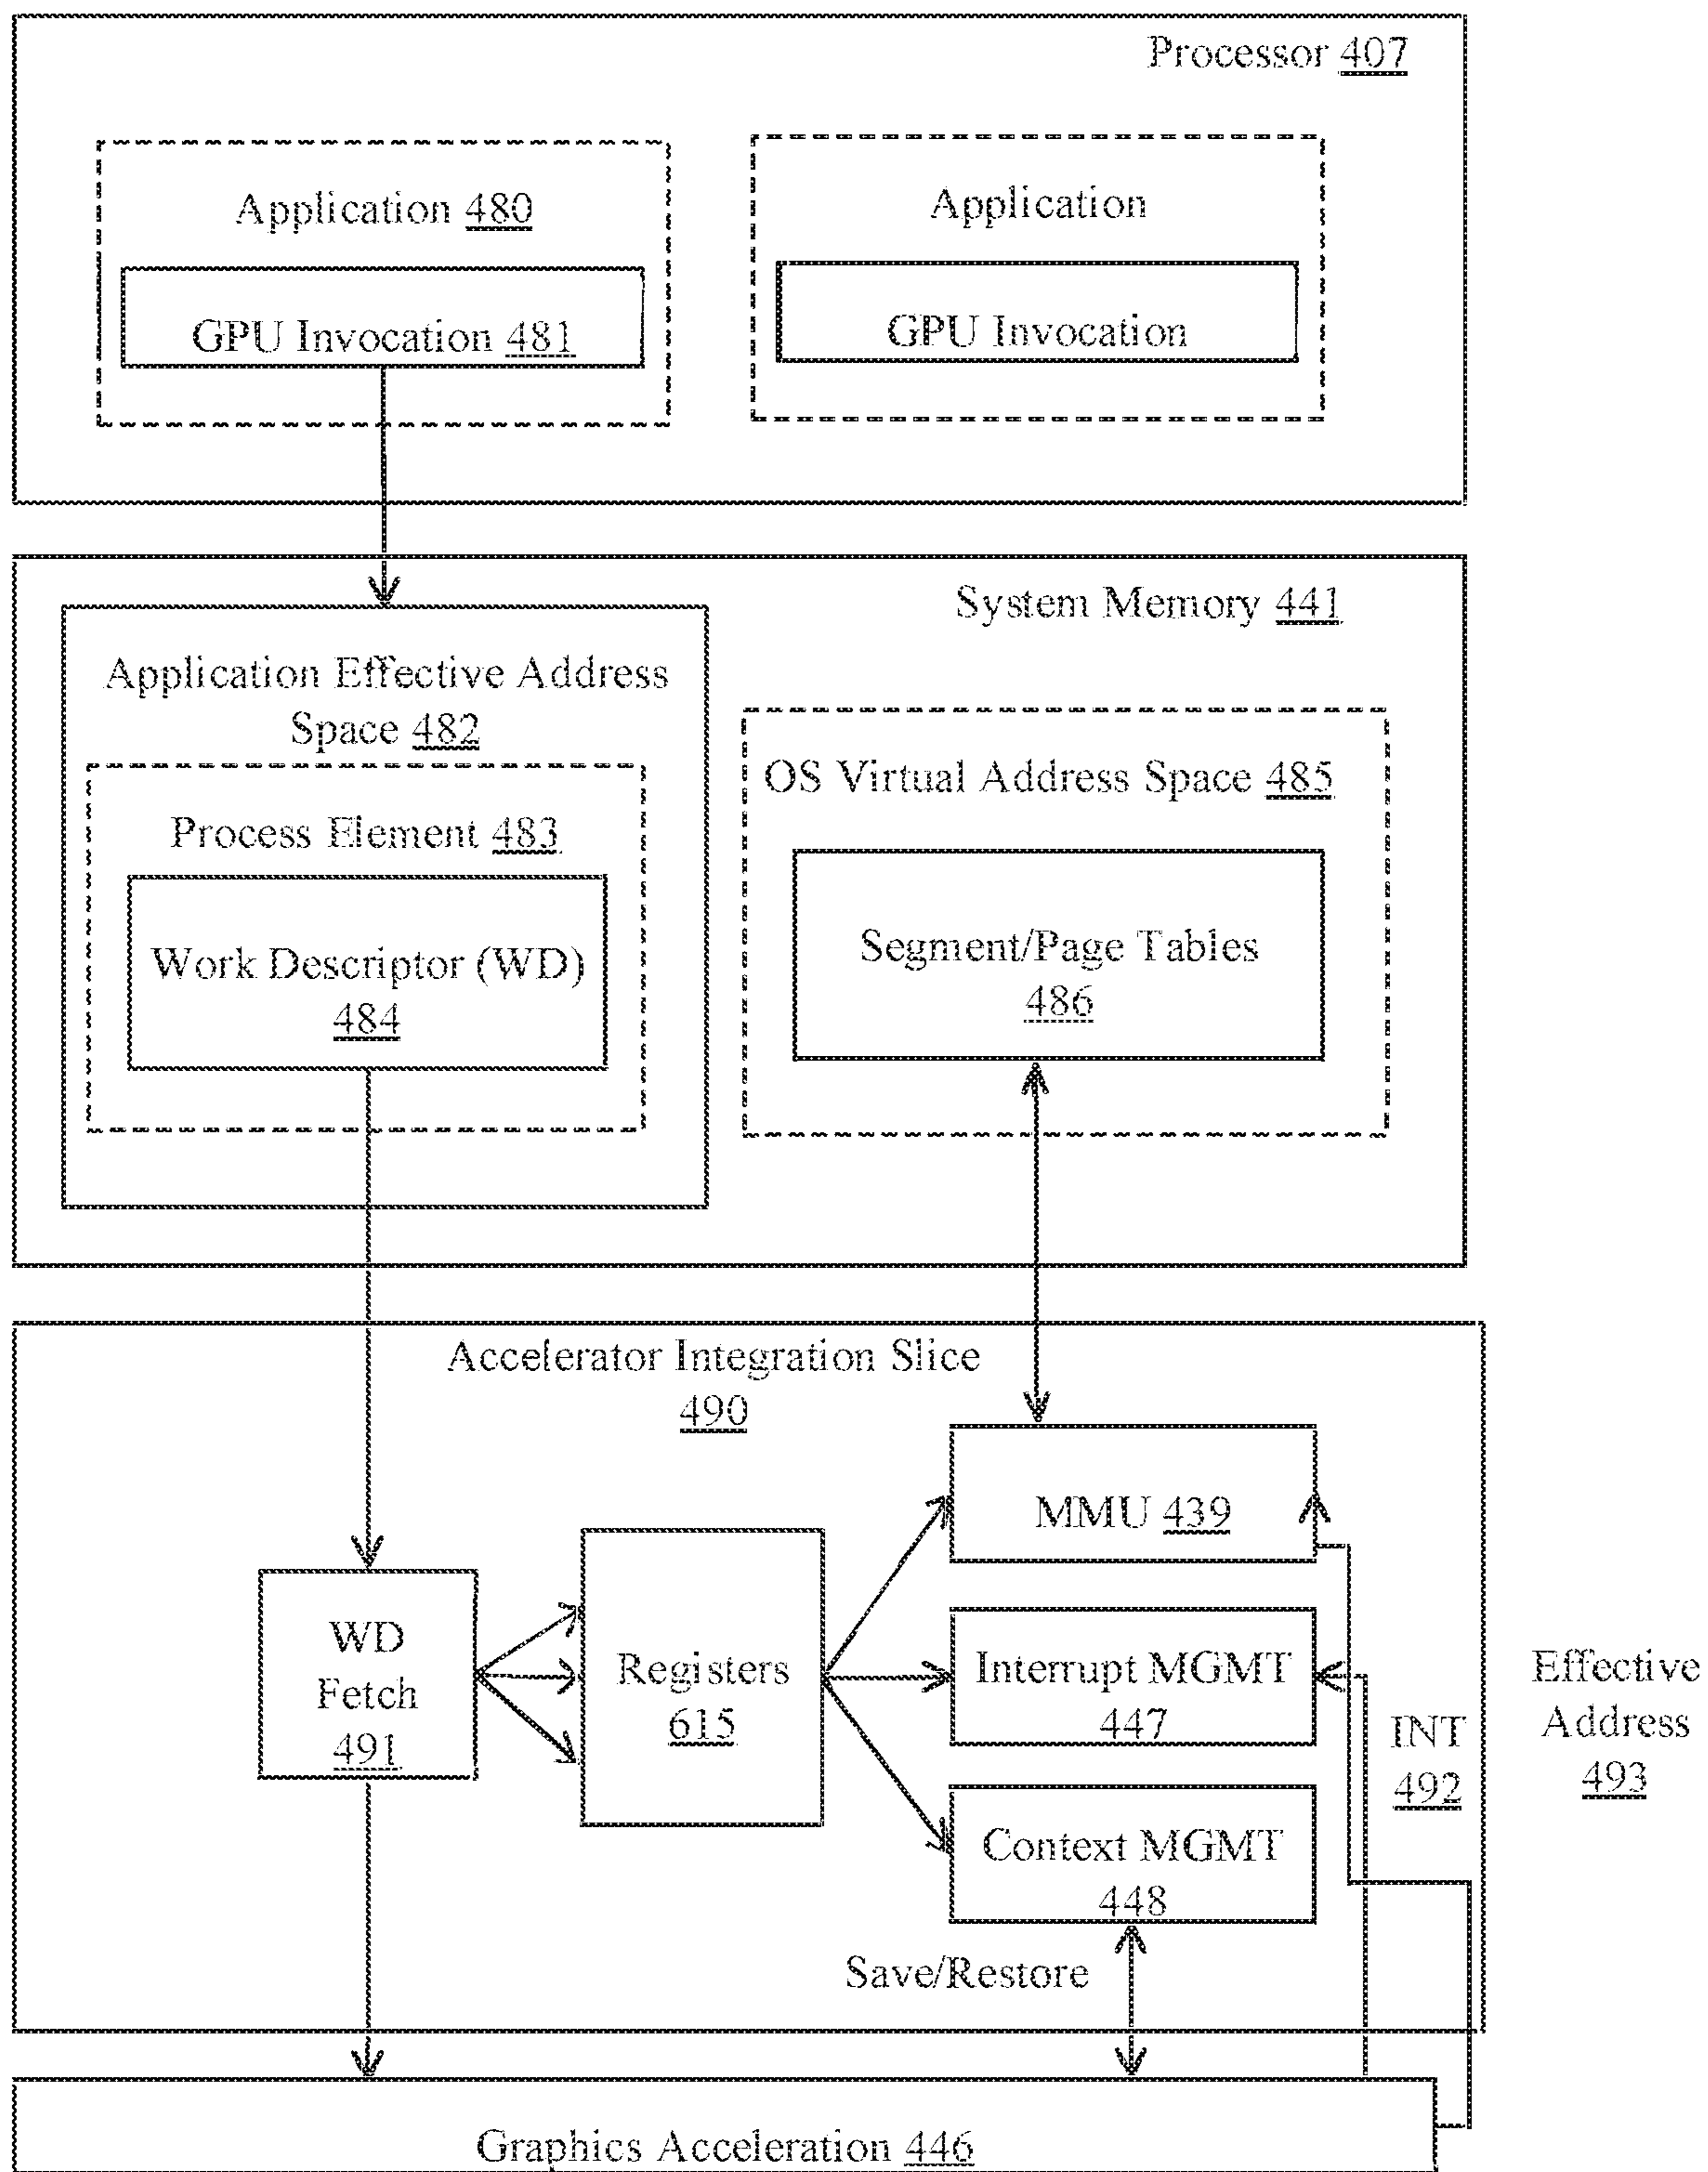


FIG. 4D

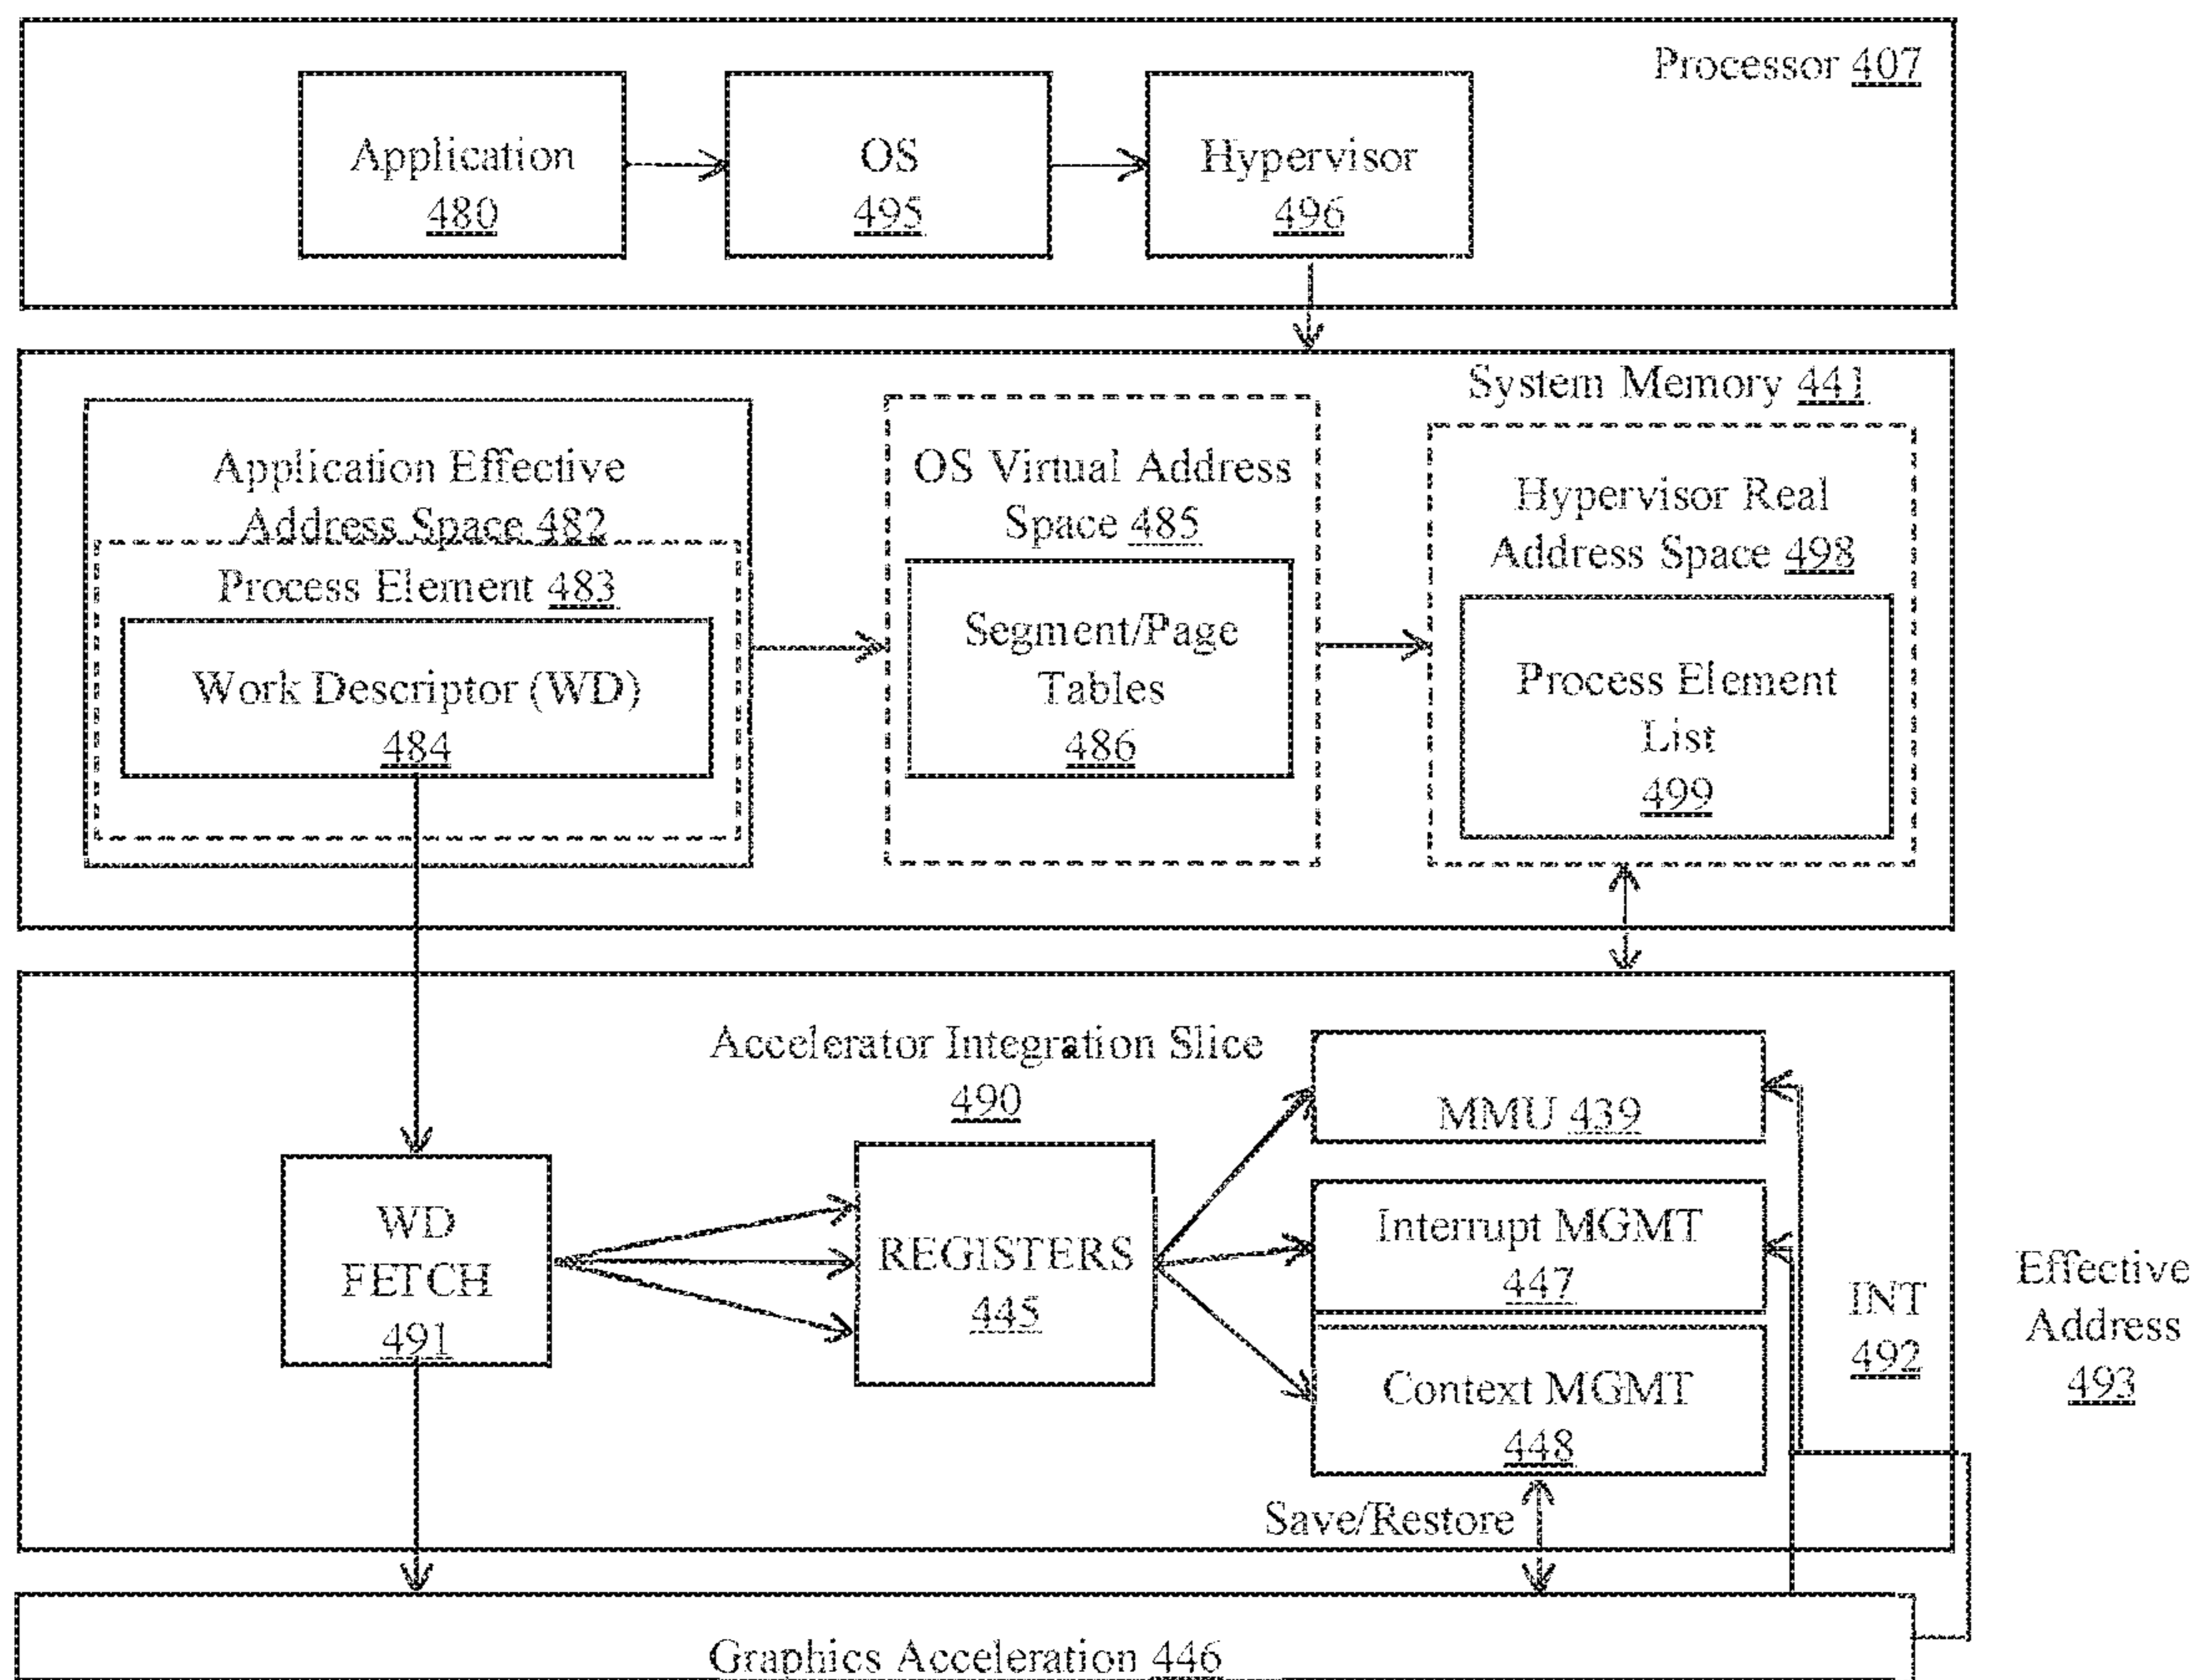


FIG. 4E

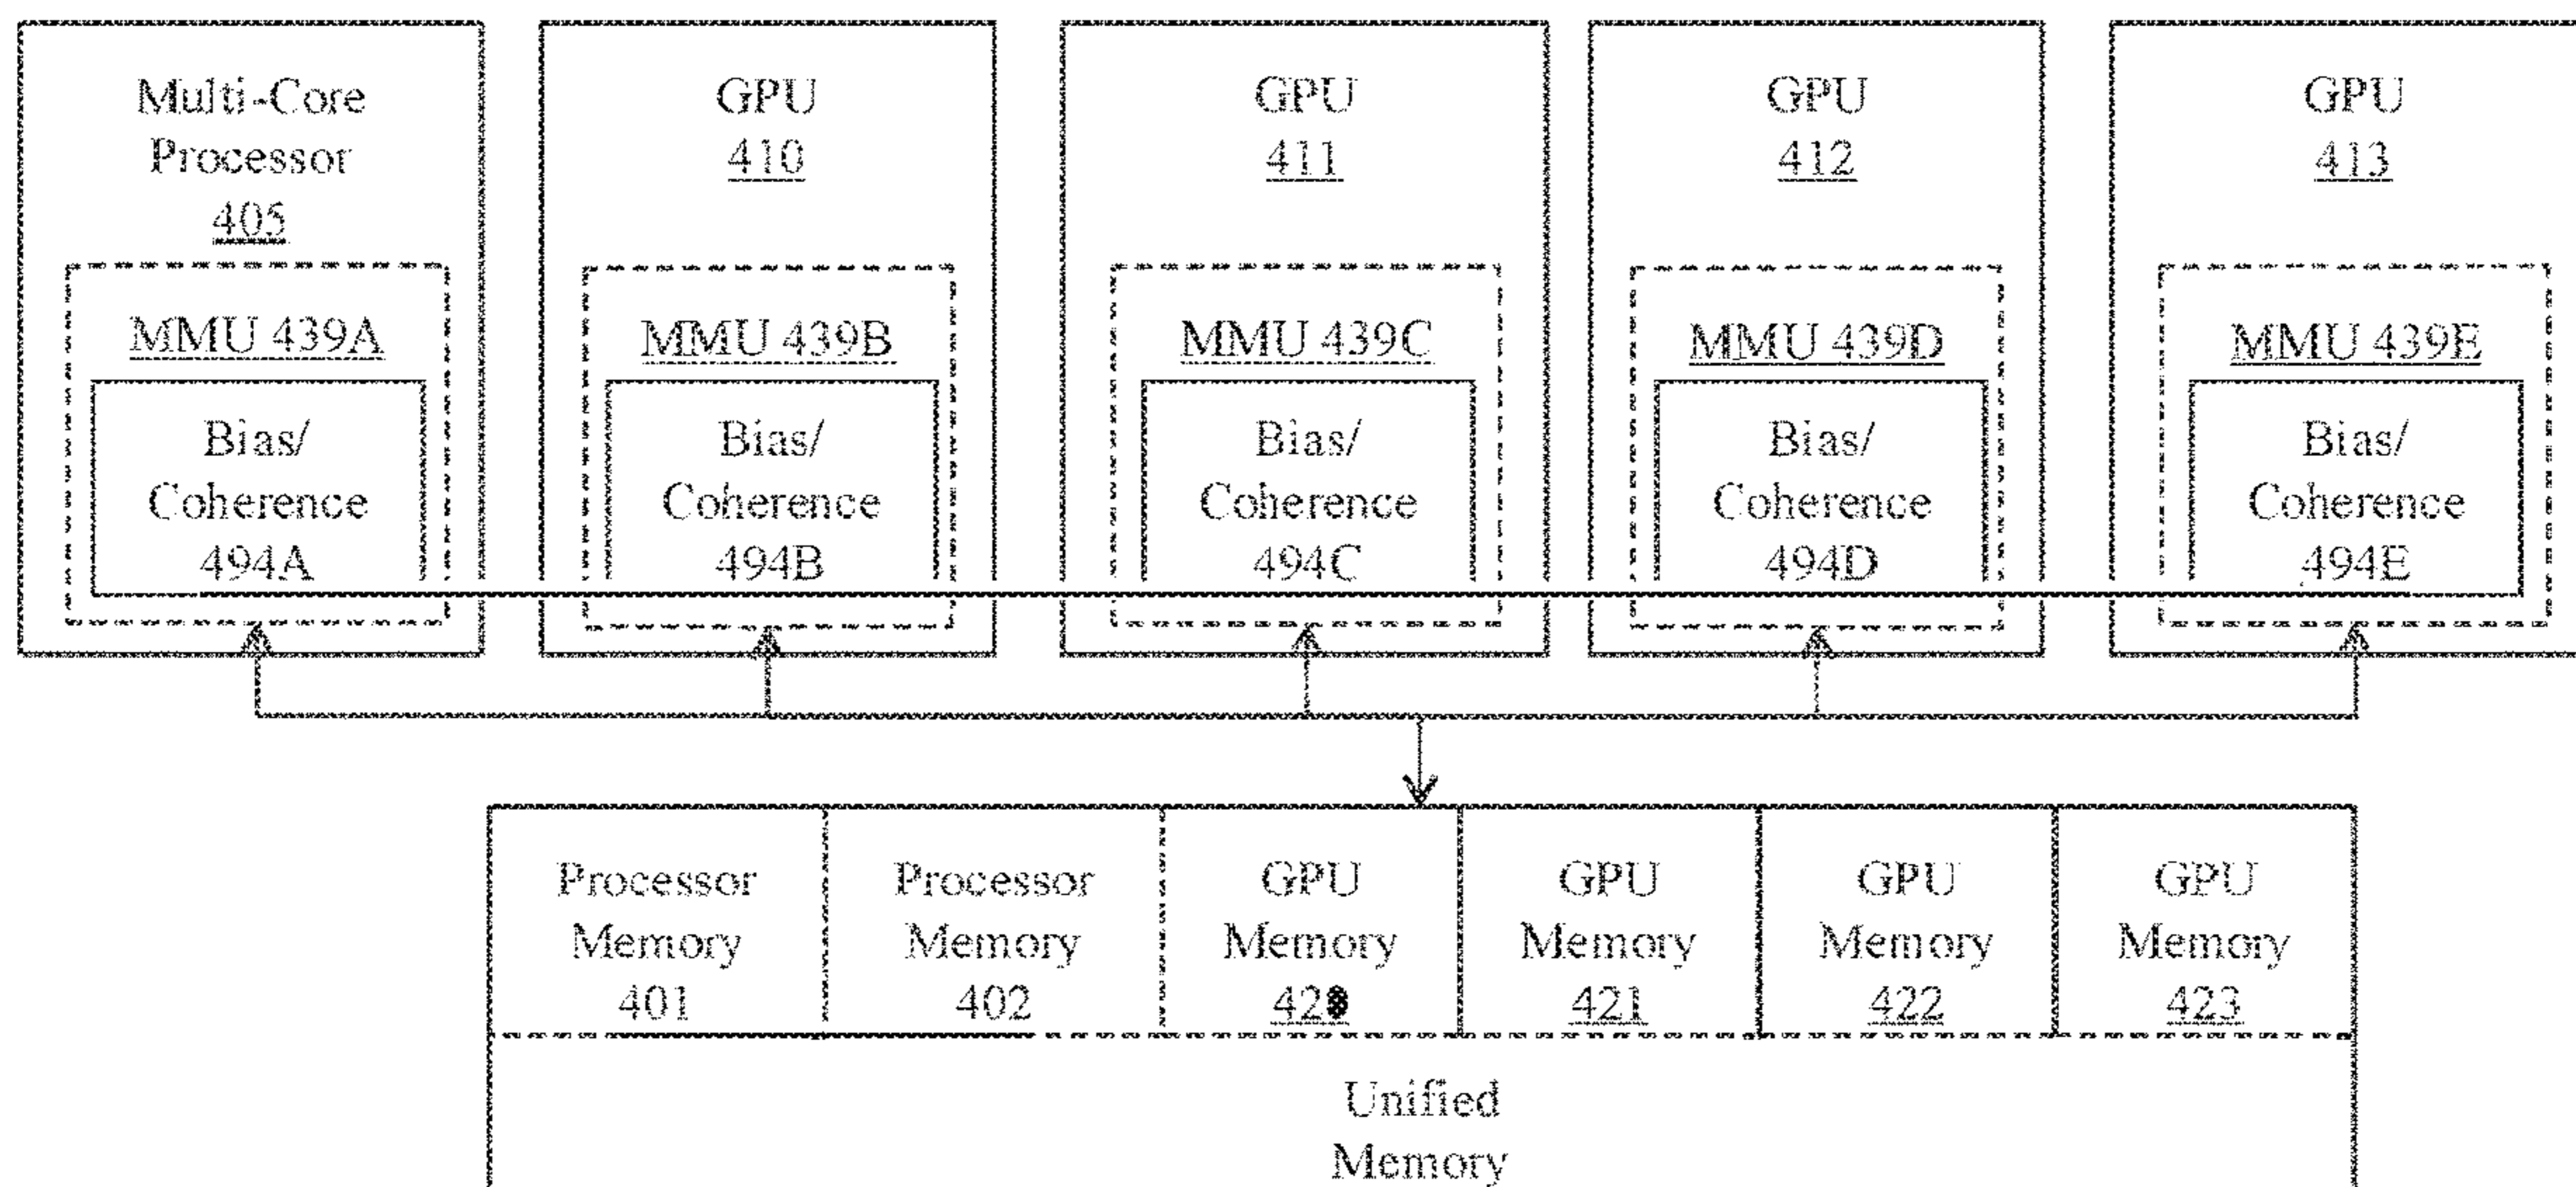


FIG. 4F

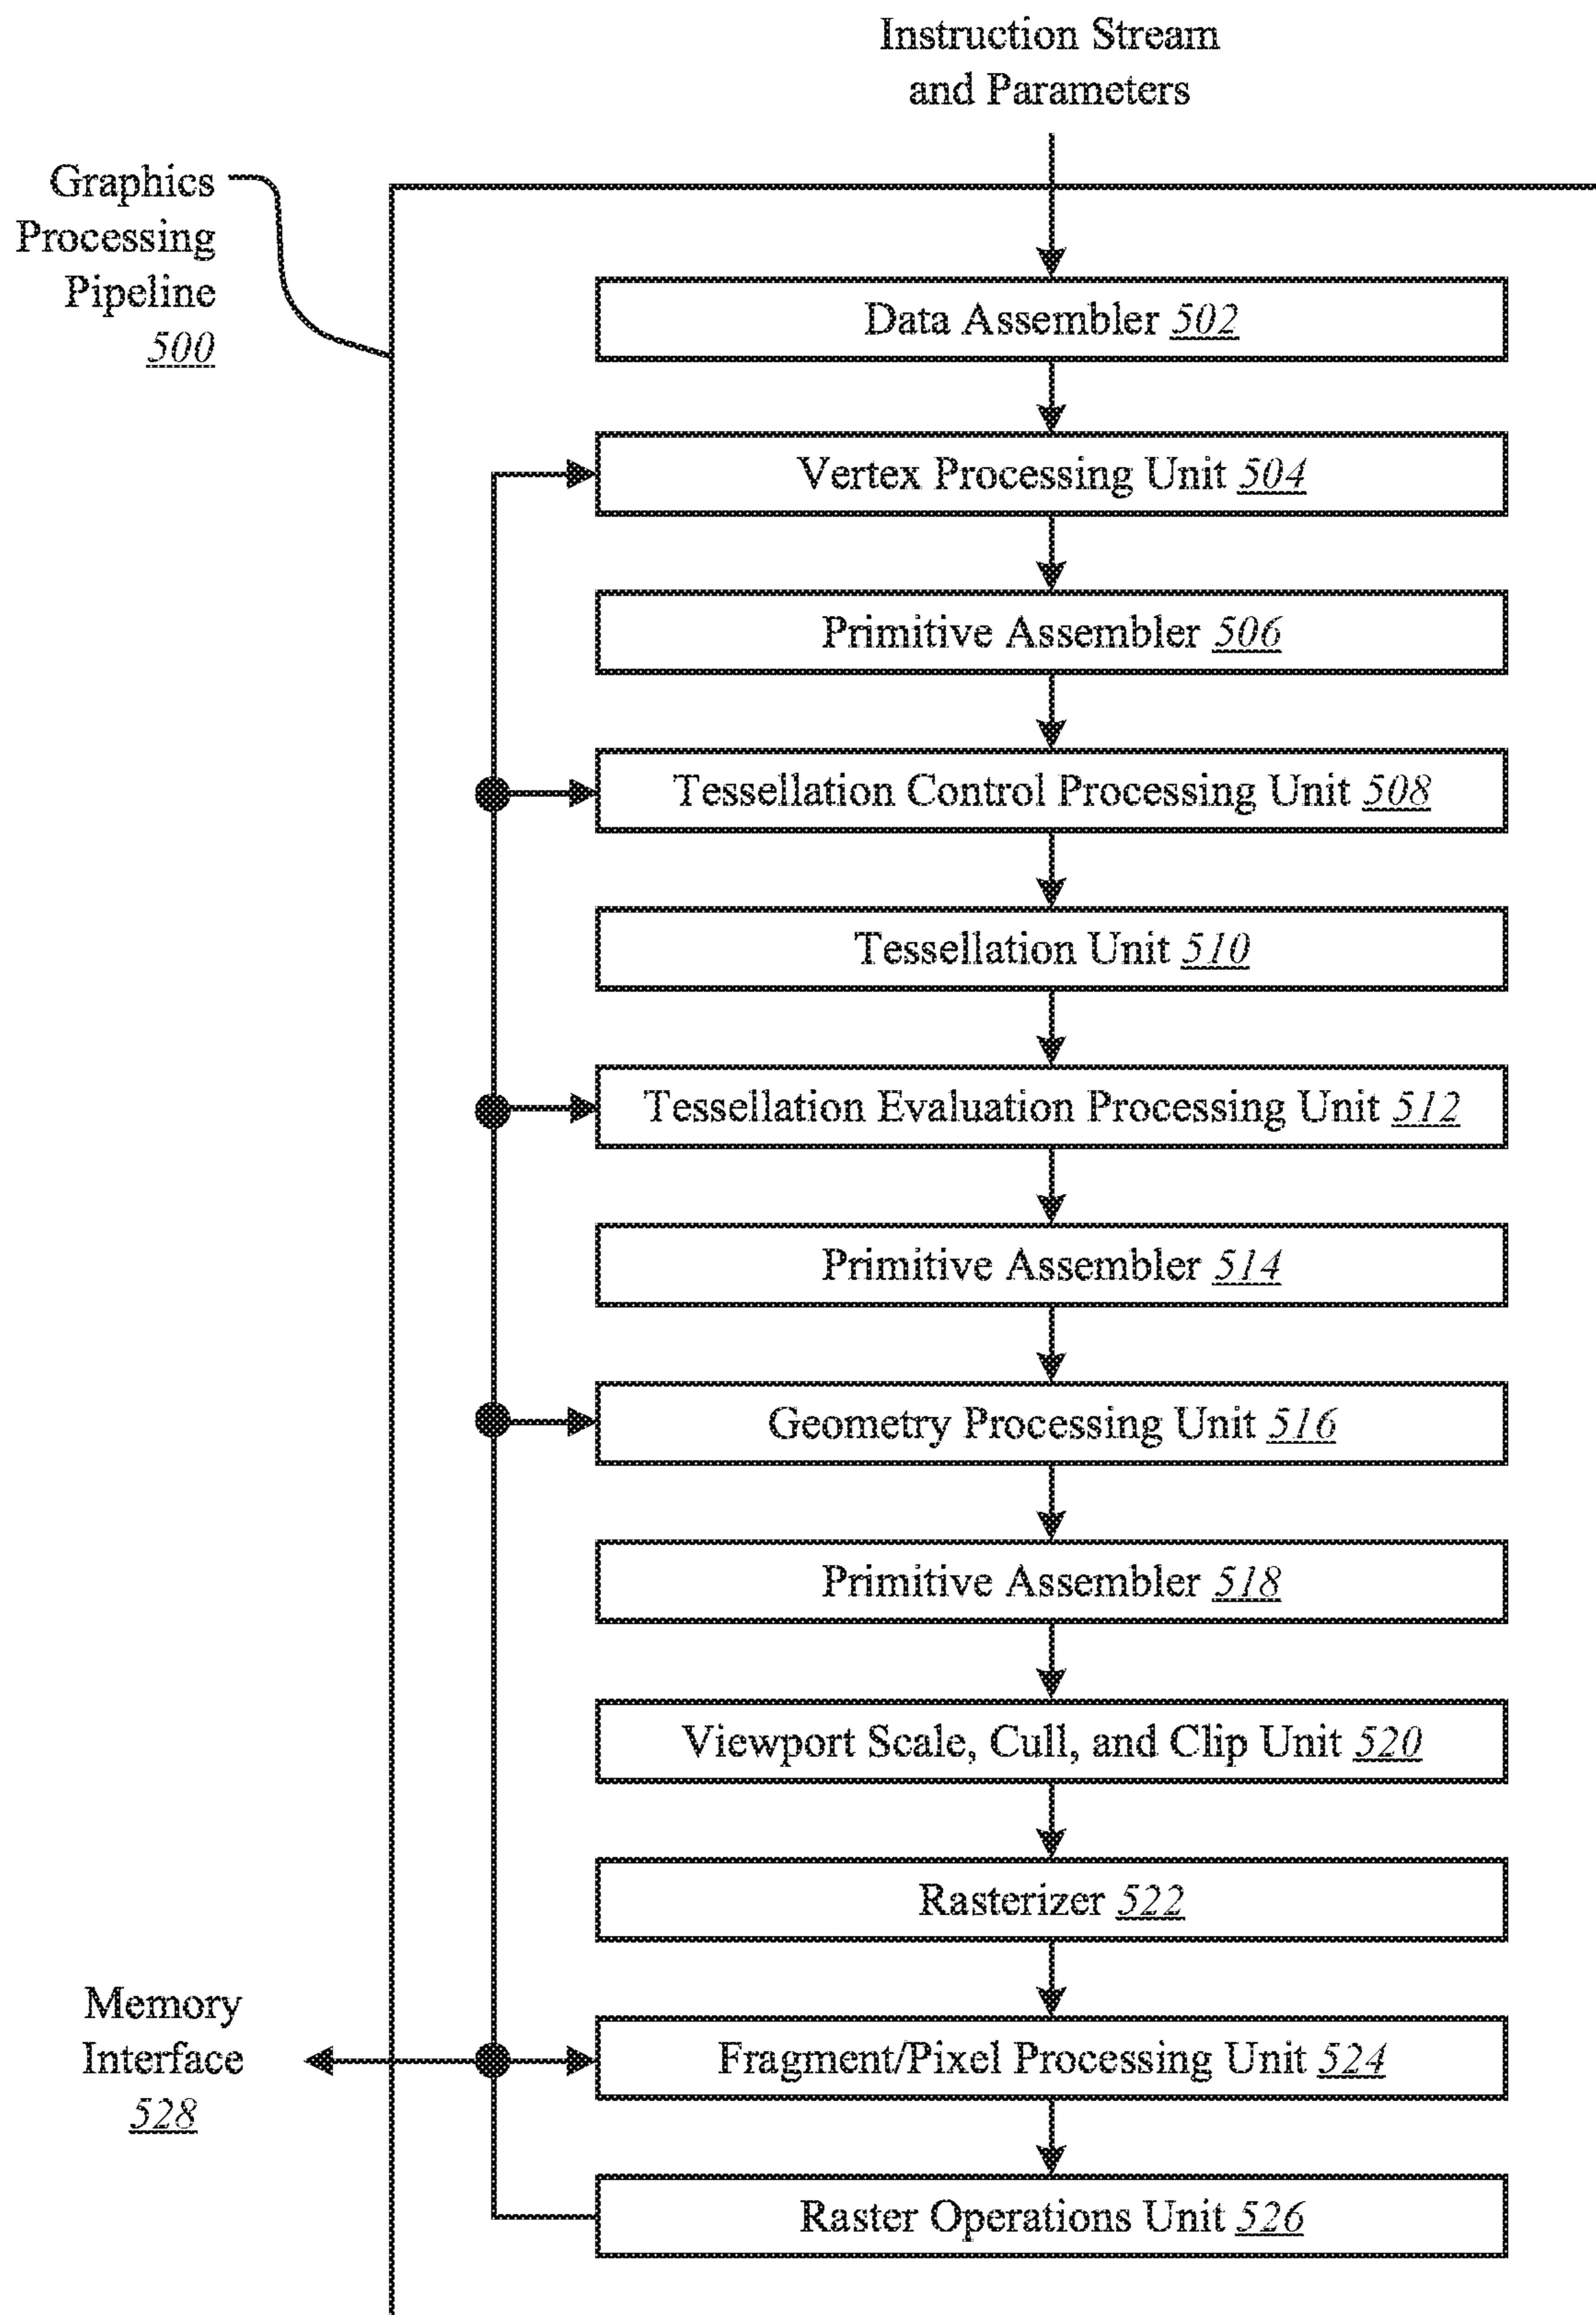
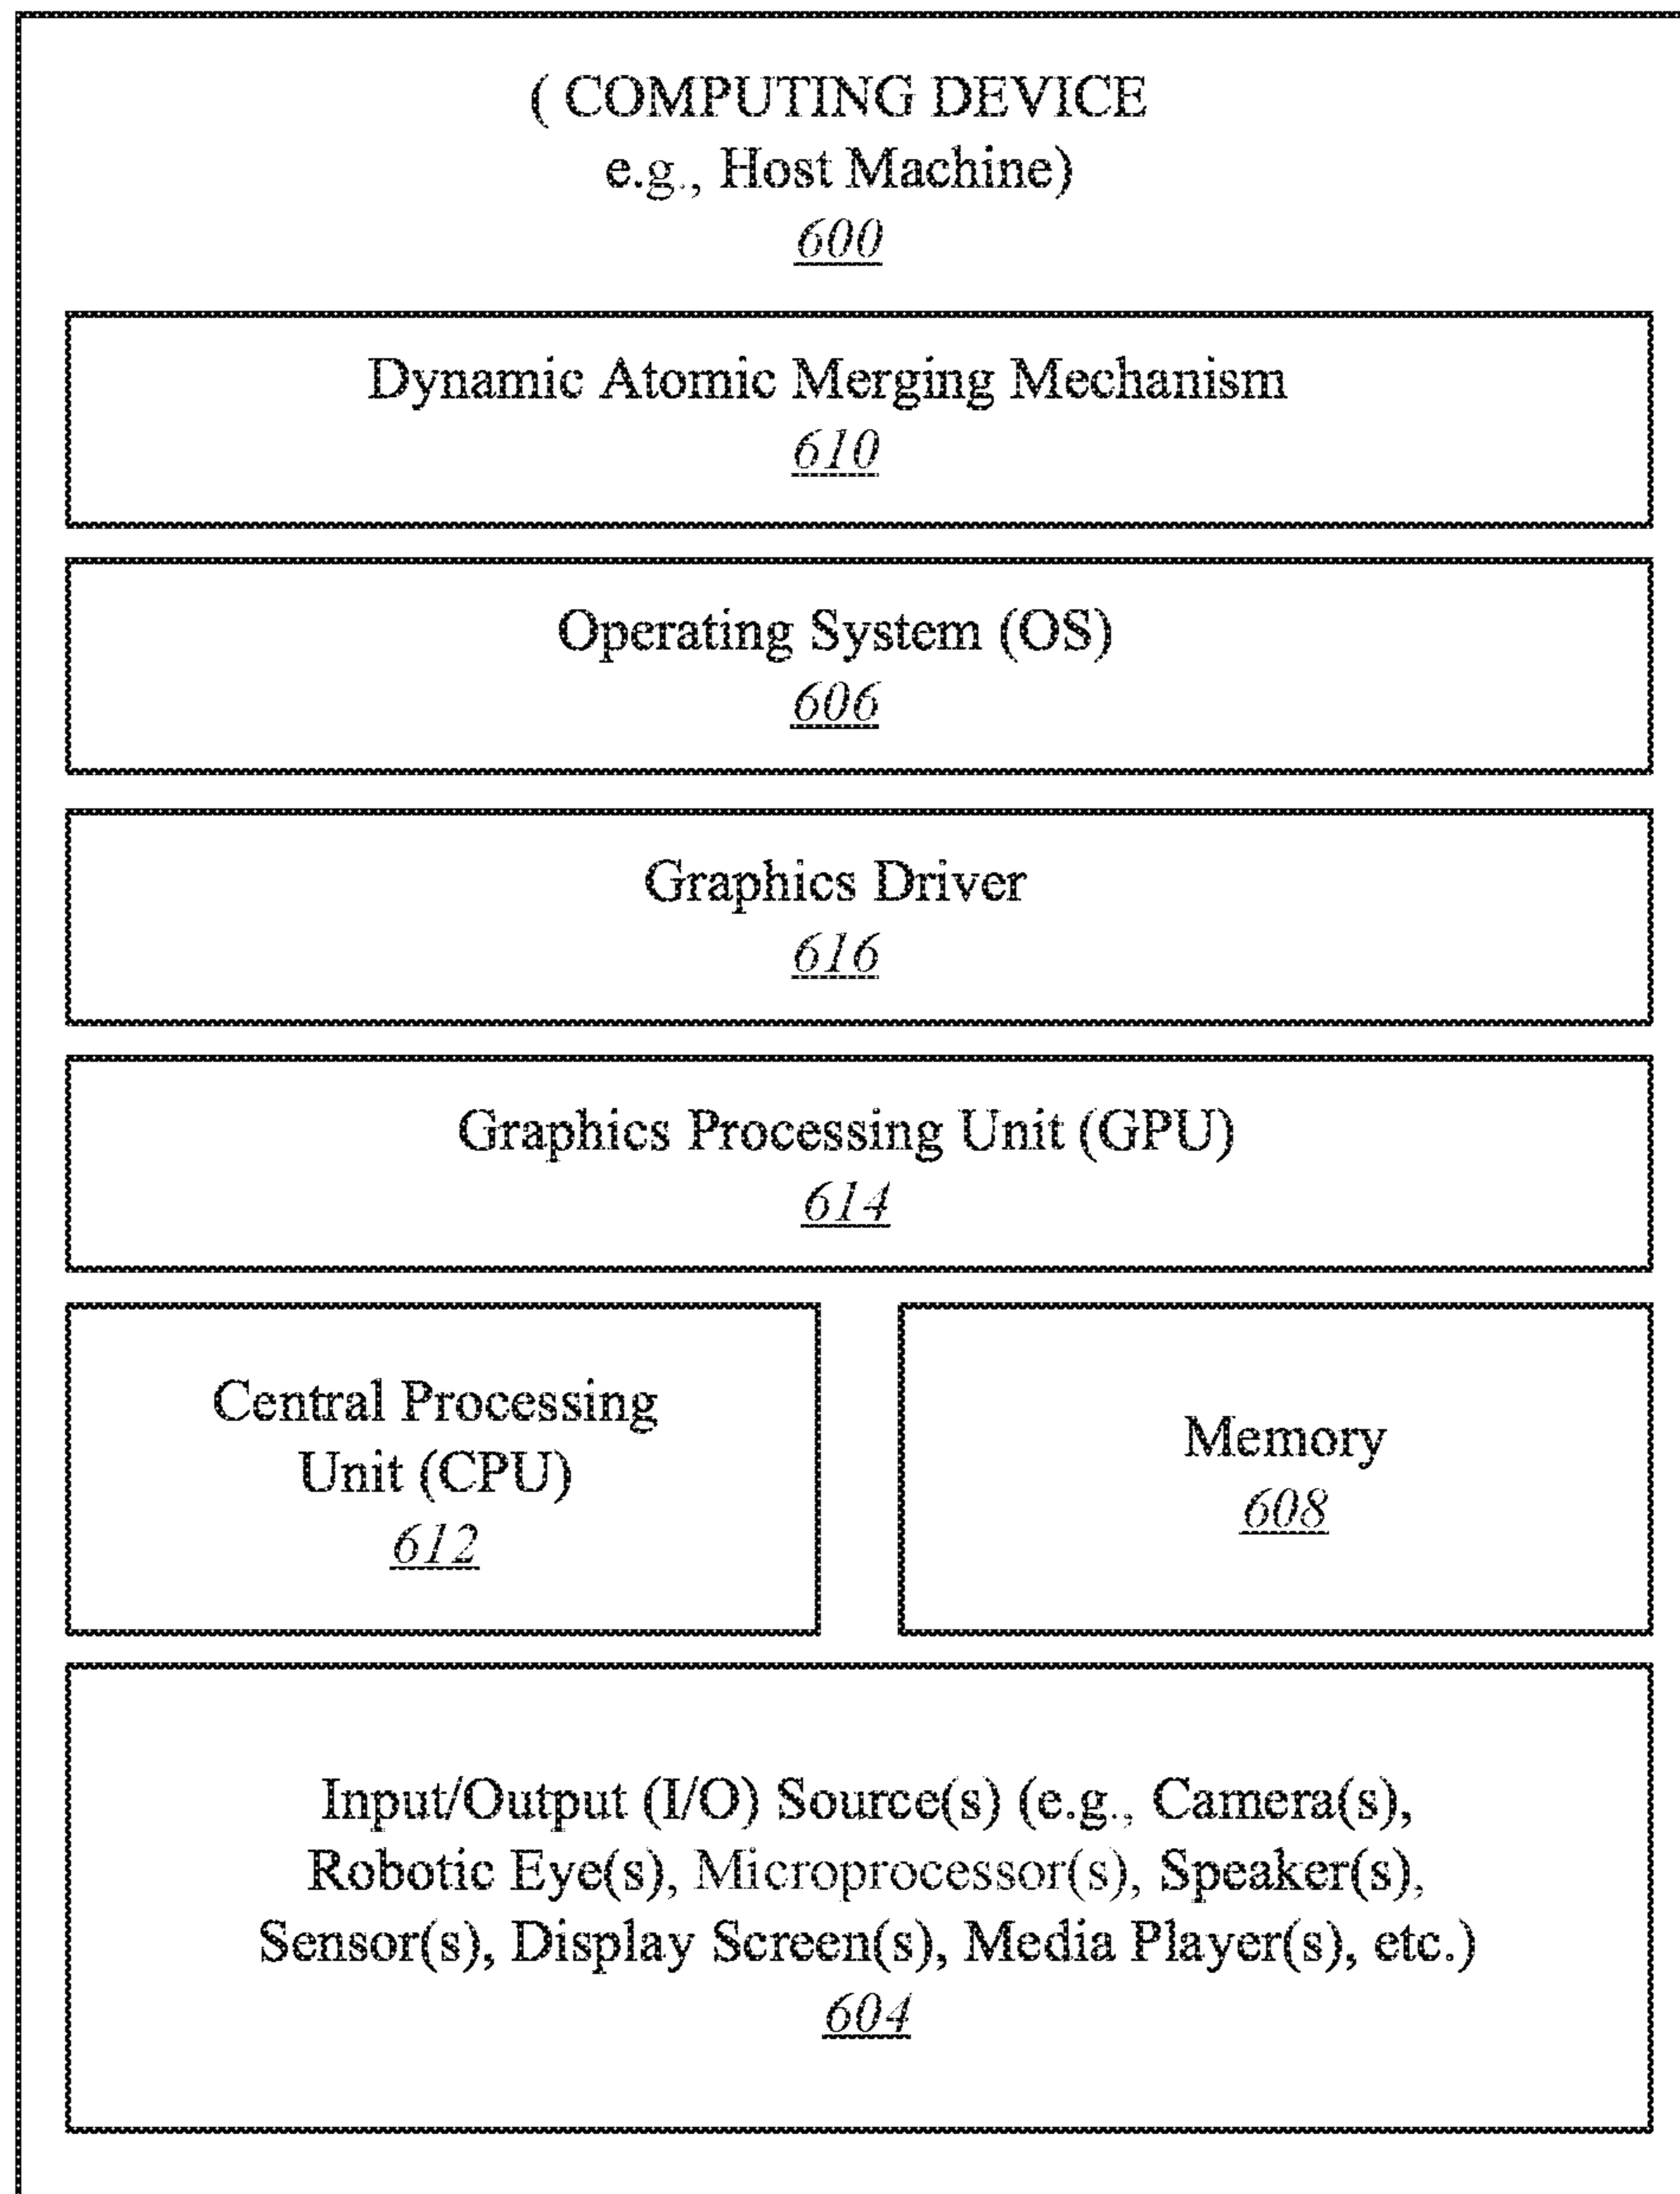


FIG. 5

*FIG. 6*


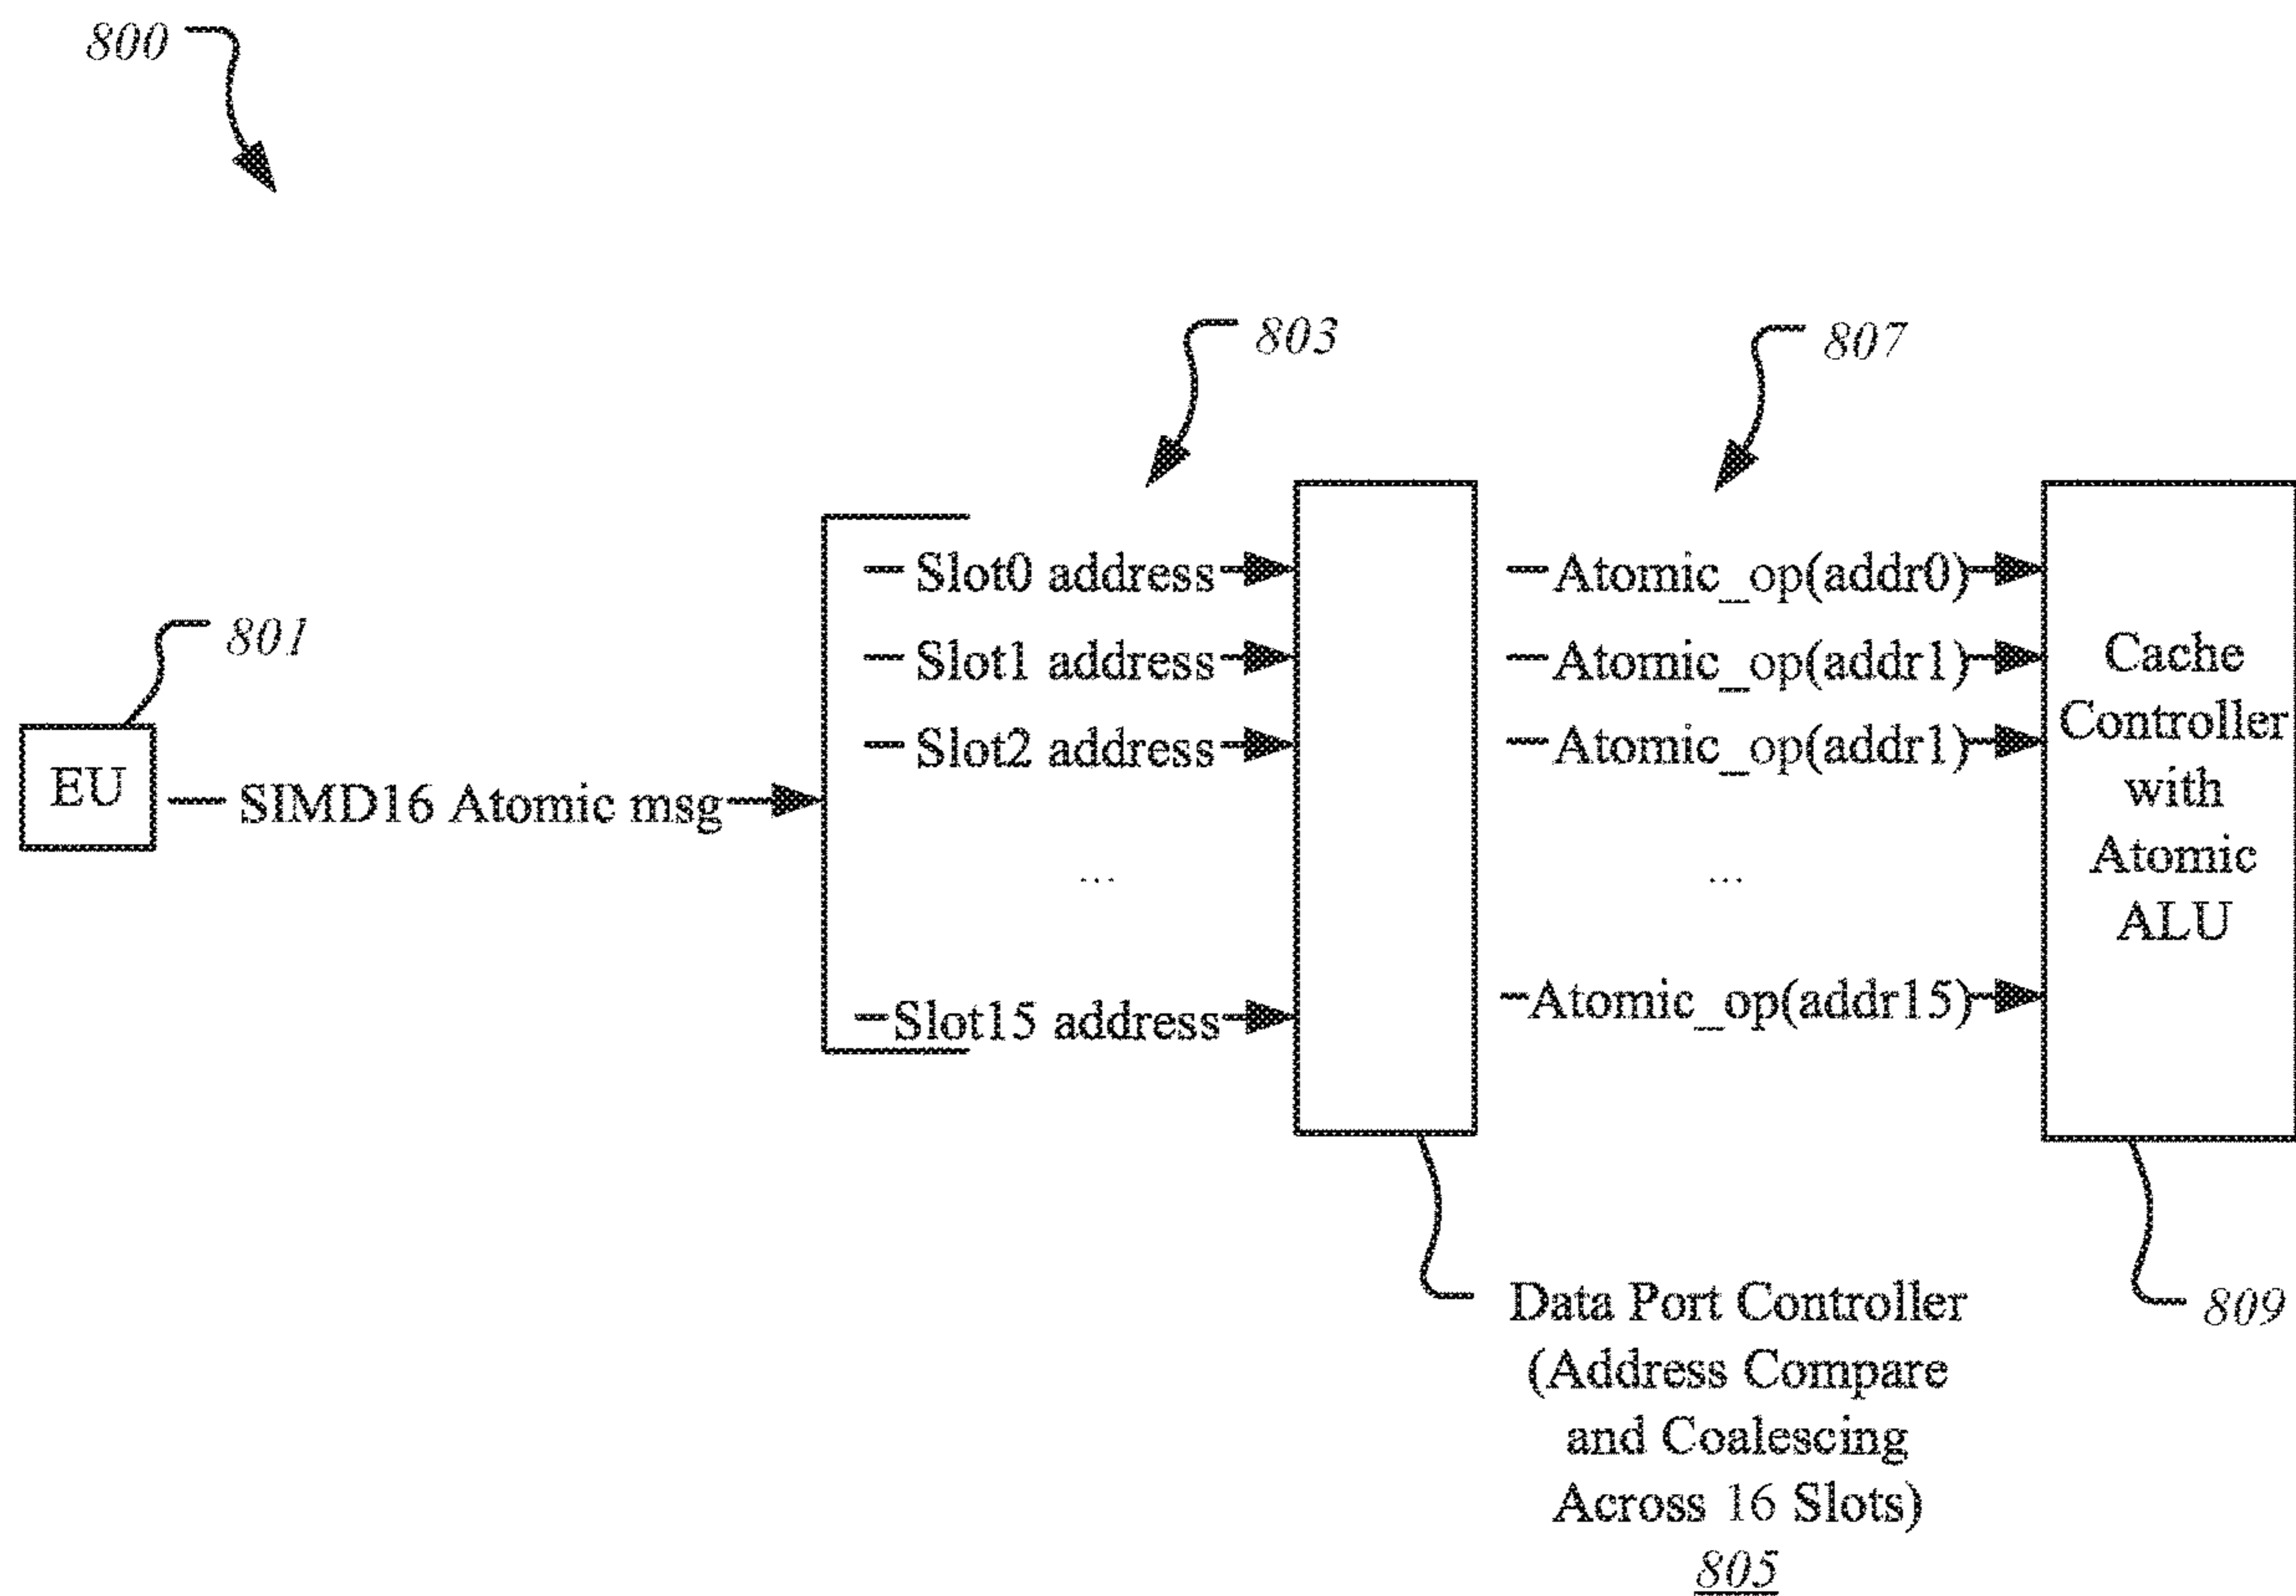


FIG. 7



PRIOR ART

FIG. 8




FIG. 9A


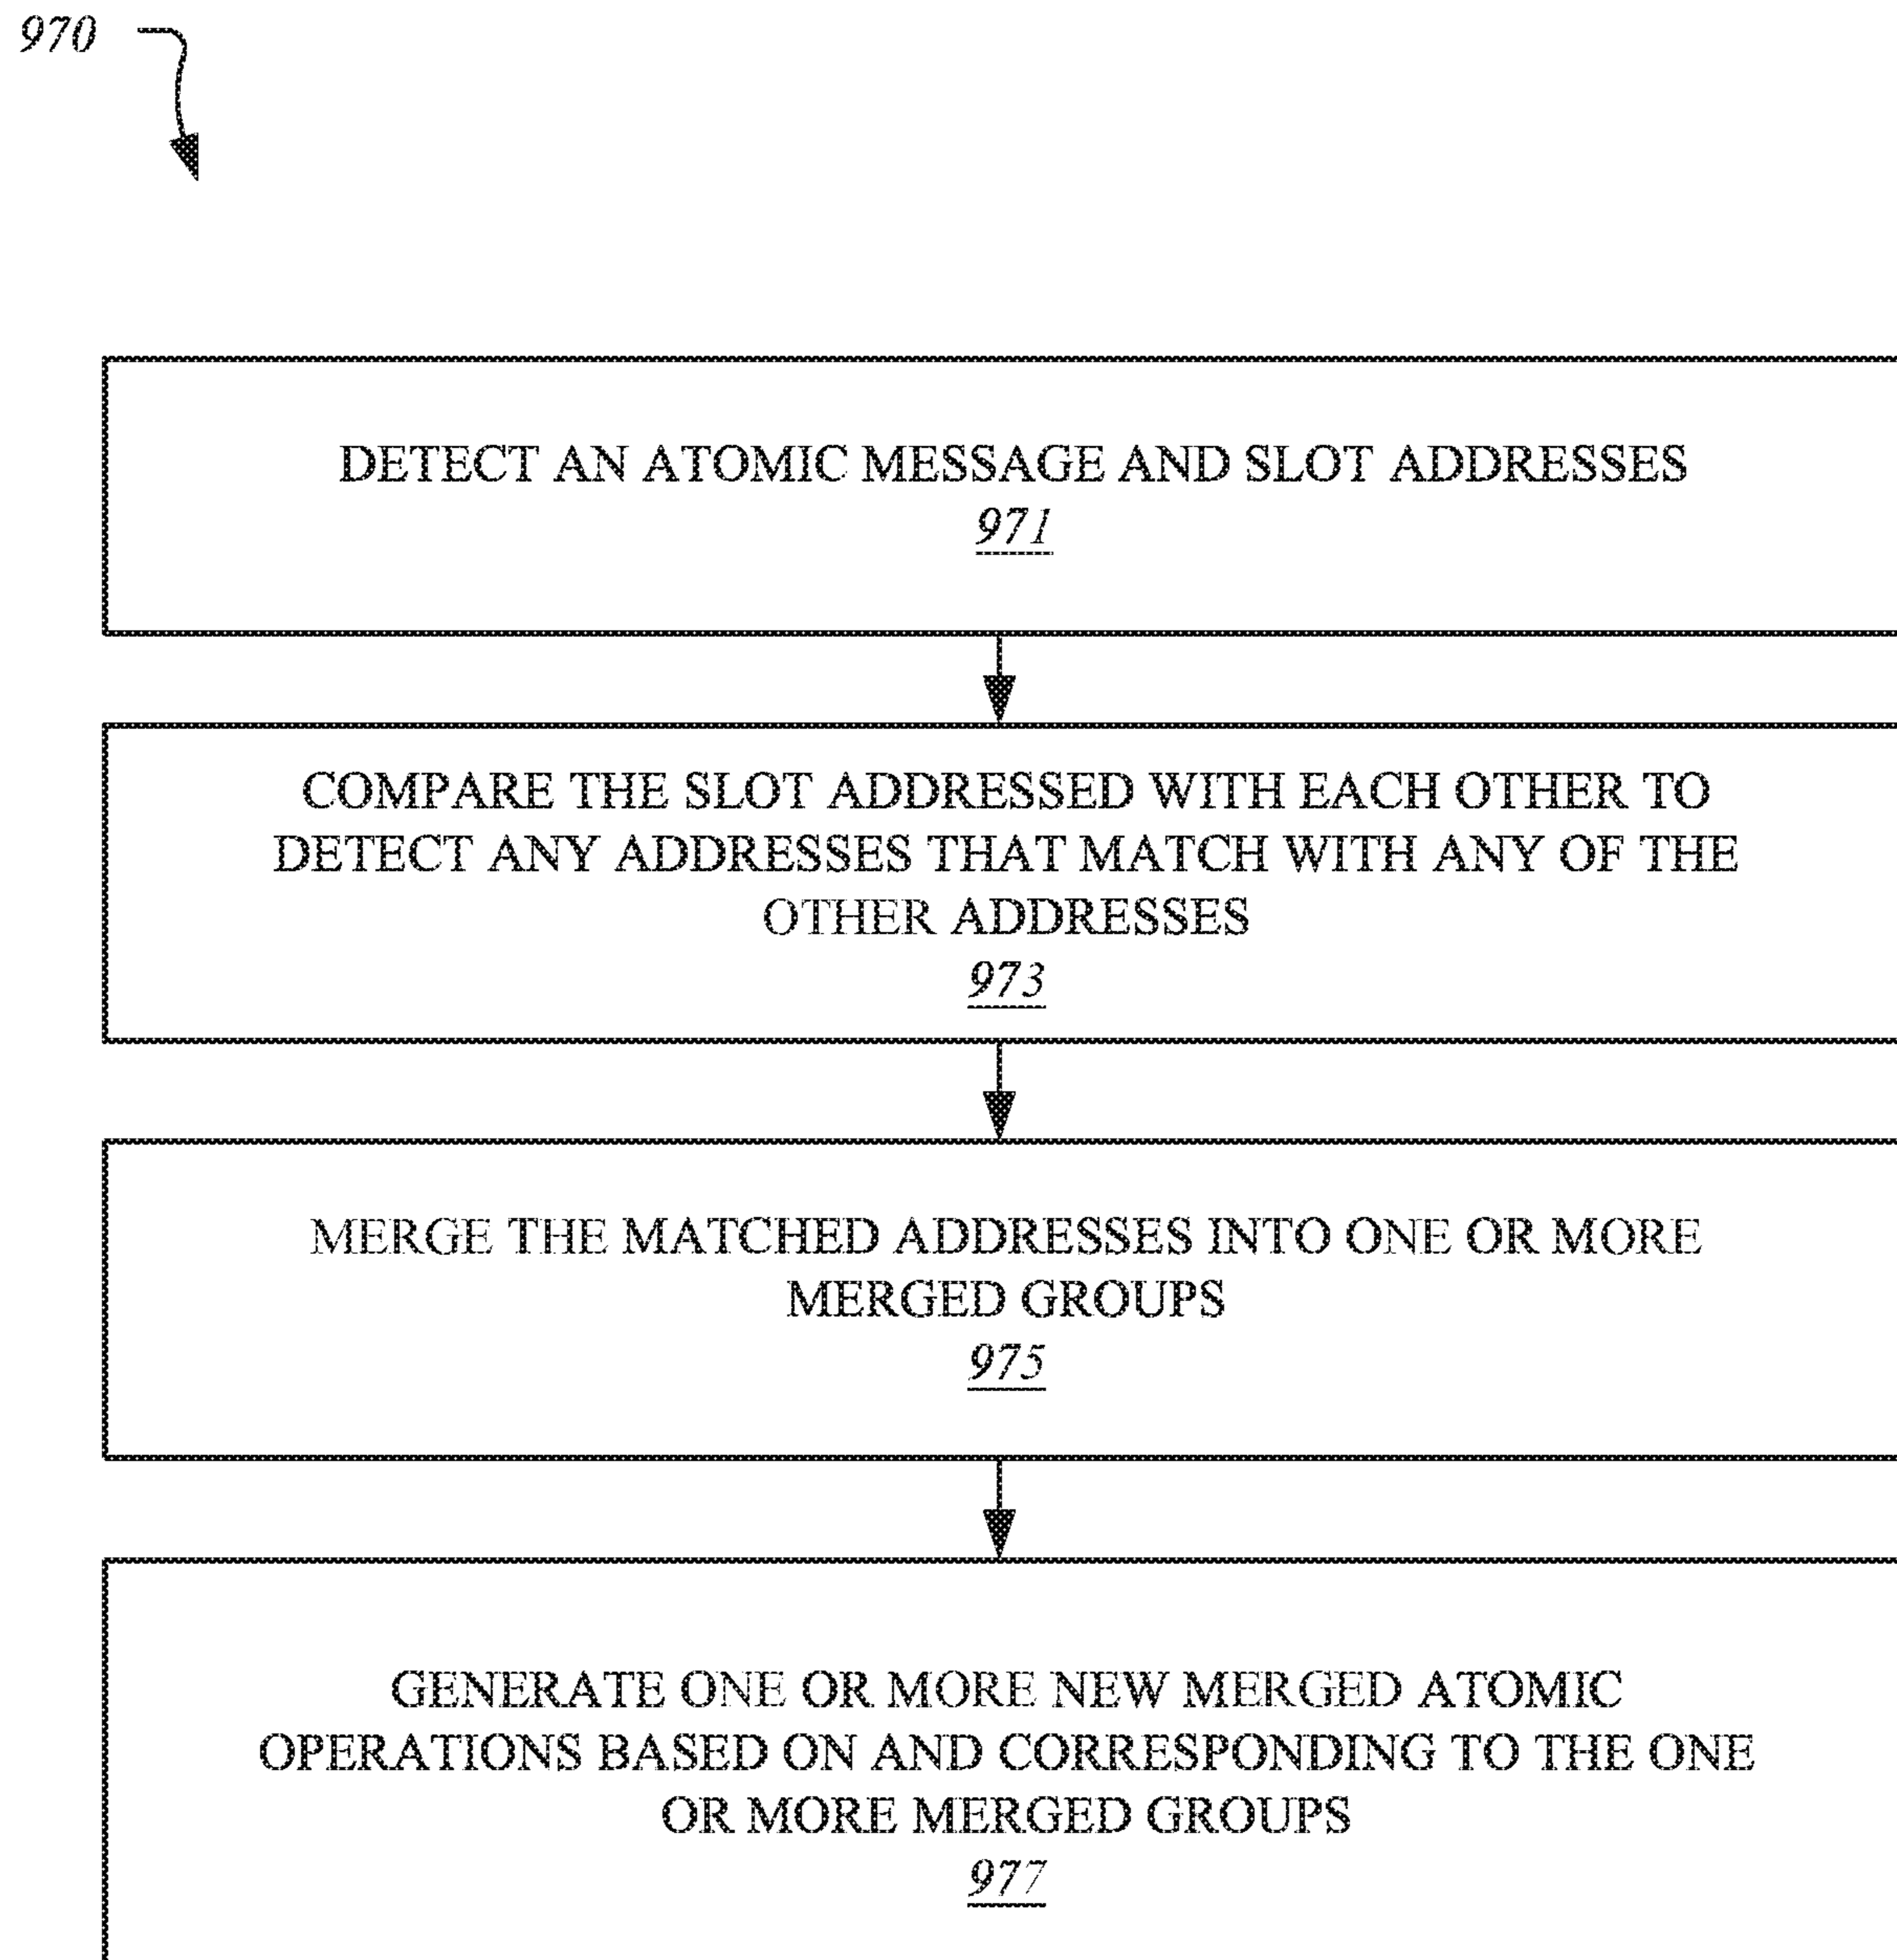


FIG. 9B

*FIG. 9C*

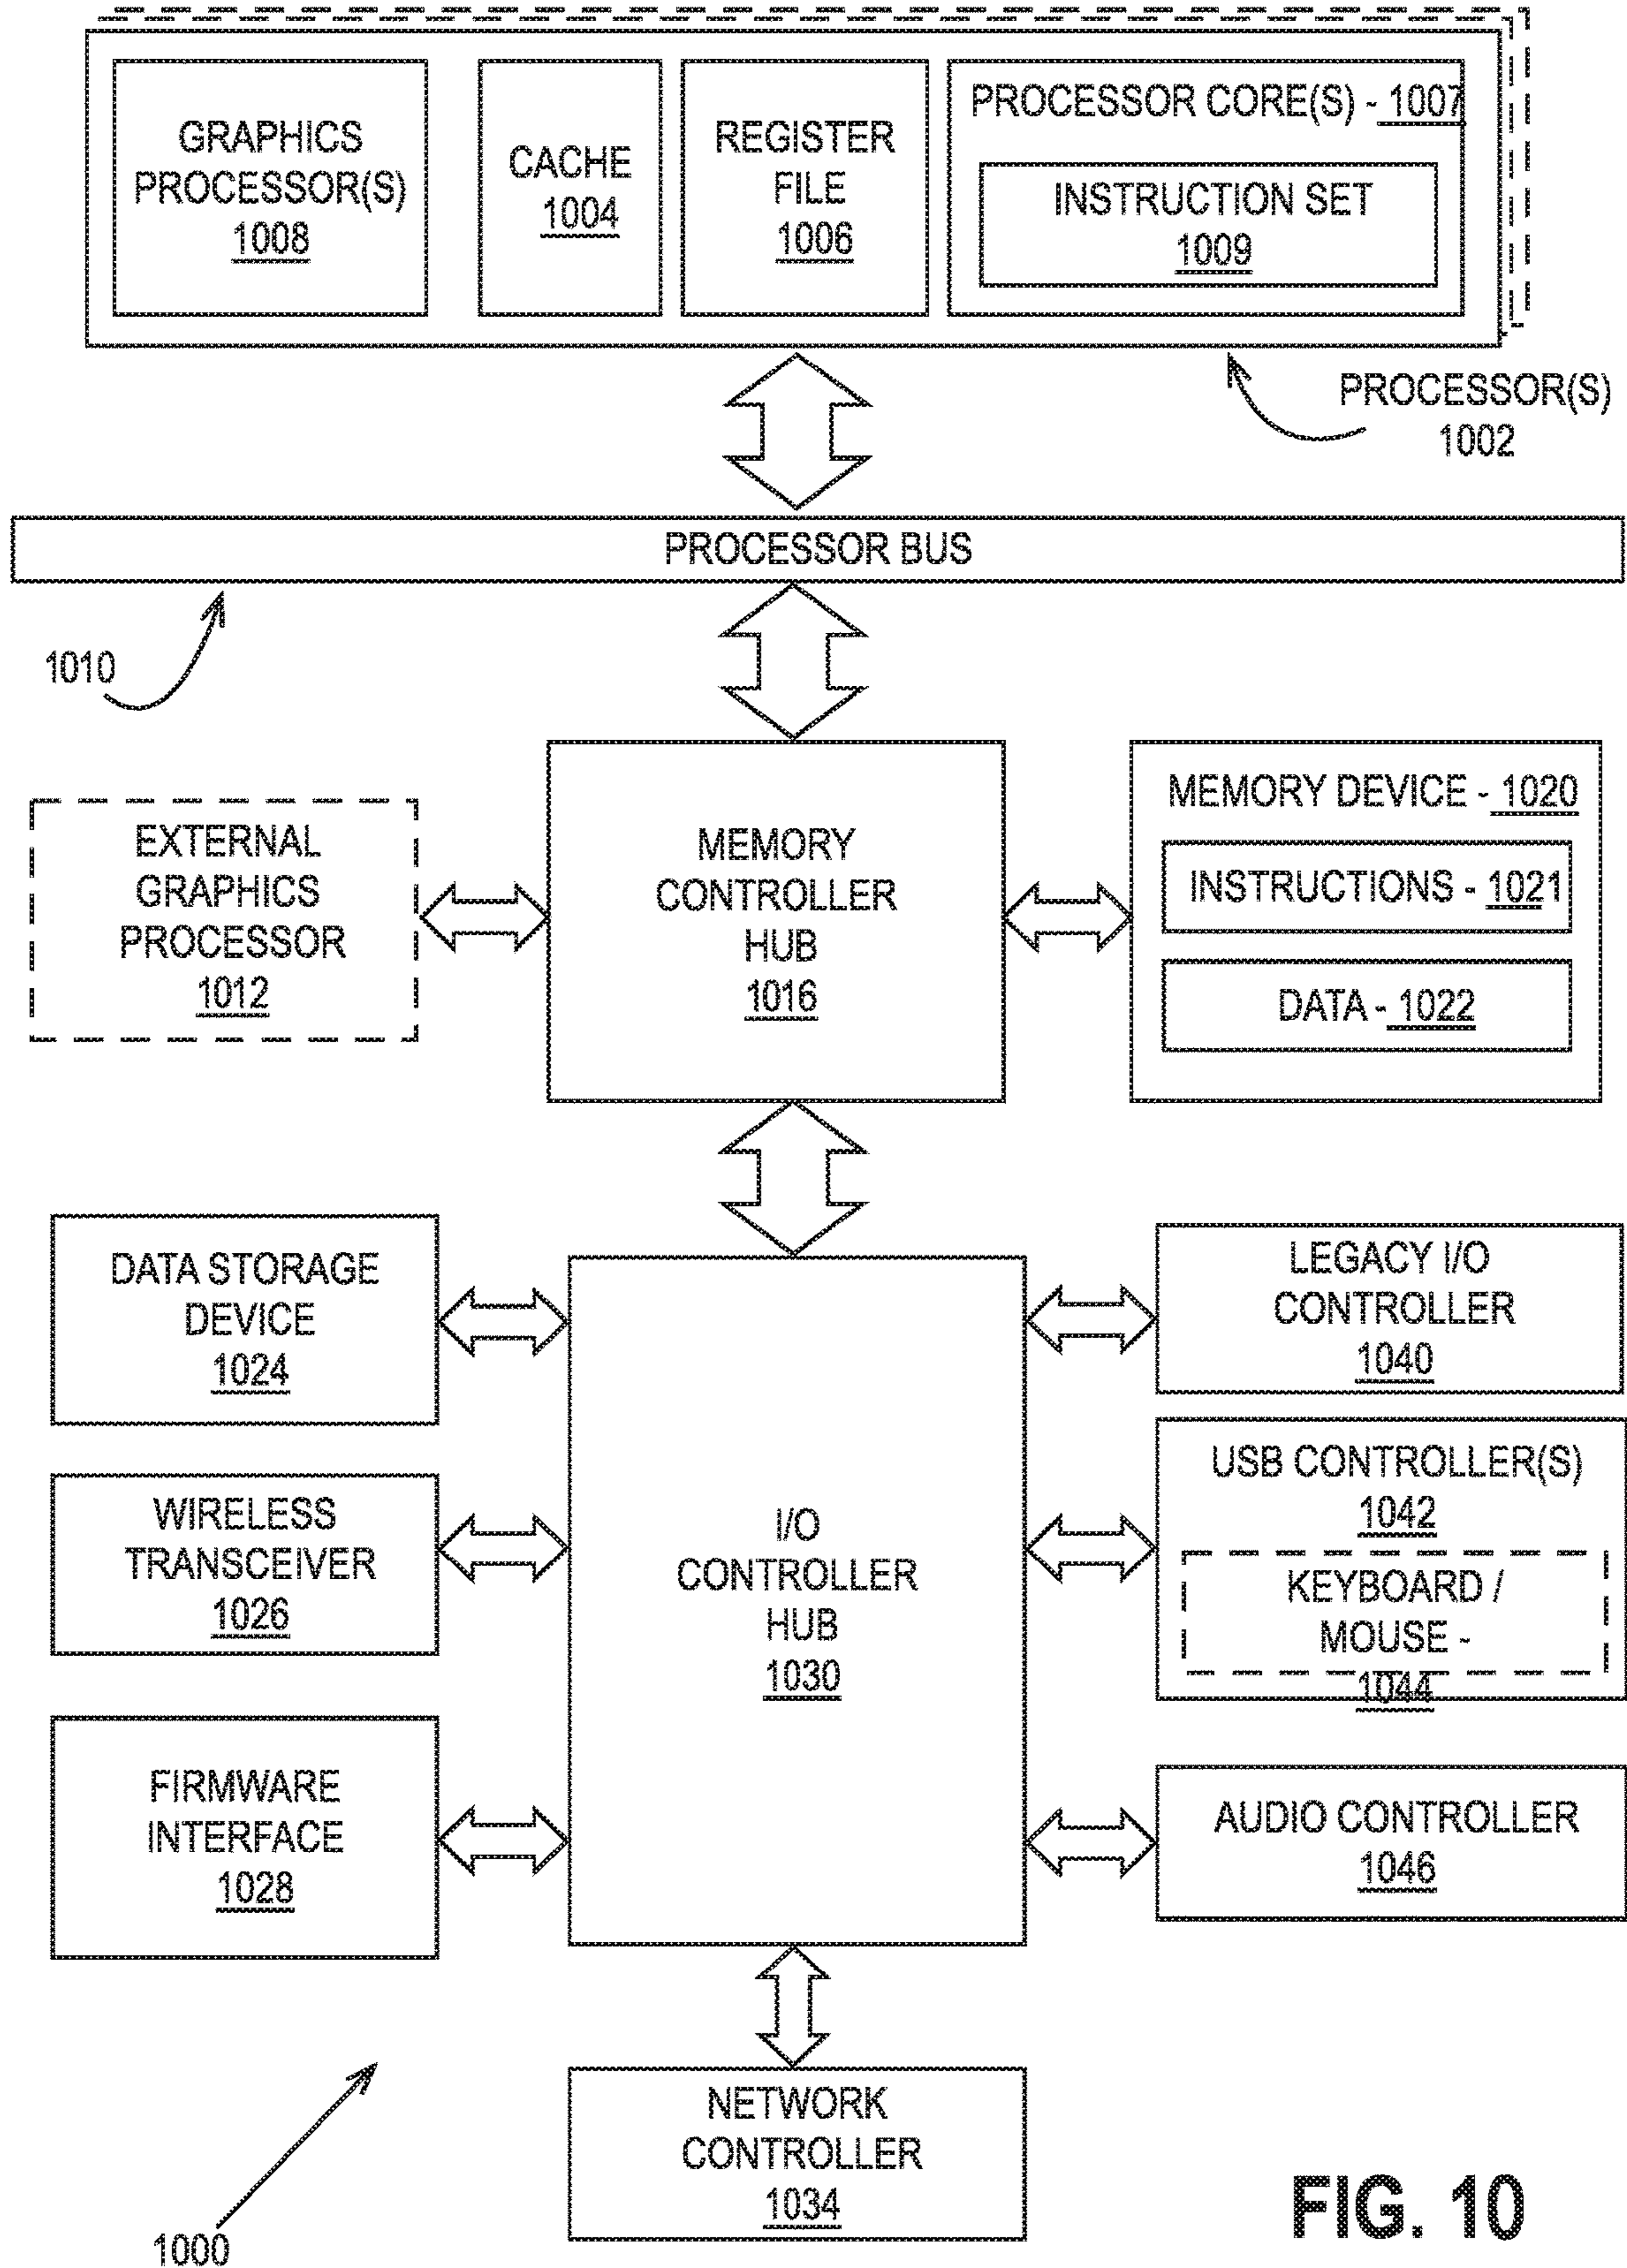


FIG. 10

PROCESSOR 1100




FIG. 11




FIG. 12




FIG. 13




FIG. 14





FIG. 15

GRAPHICS PROCESSOR INSTRUCTION FORMATS

1600


128-BIT INSTRUCTION

1610



64-BIT COMPACT INSTRUCTION

1630



OPCODE DECODE

1640




FIG. 16




FIG. 17

FIG. 18A

GRAPHICS PROCESSOR COMMAND FORMAT

1800


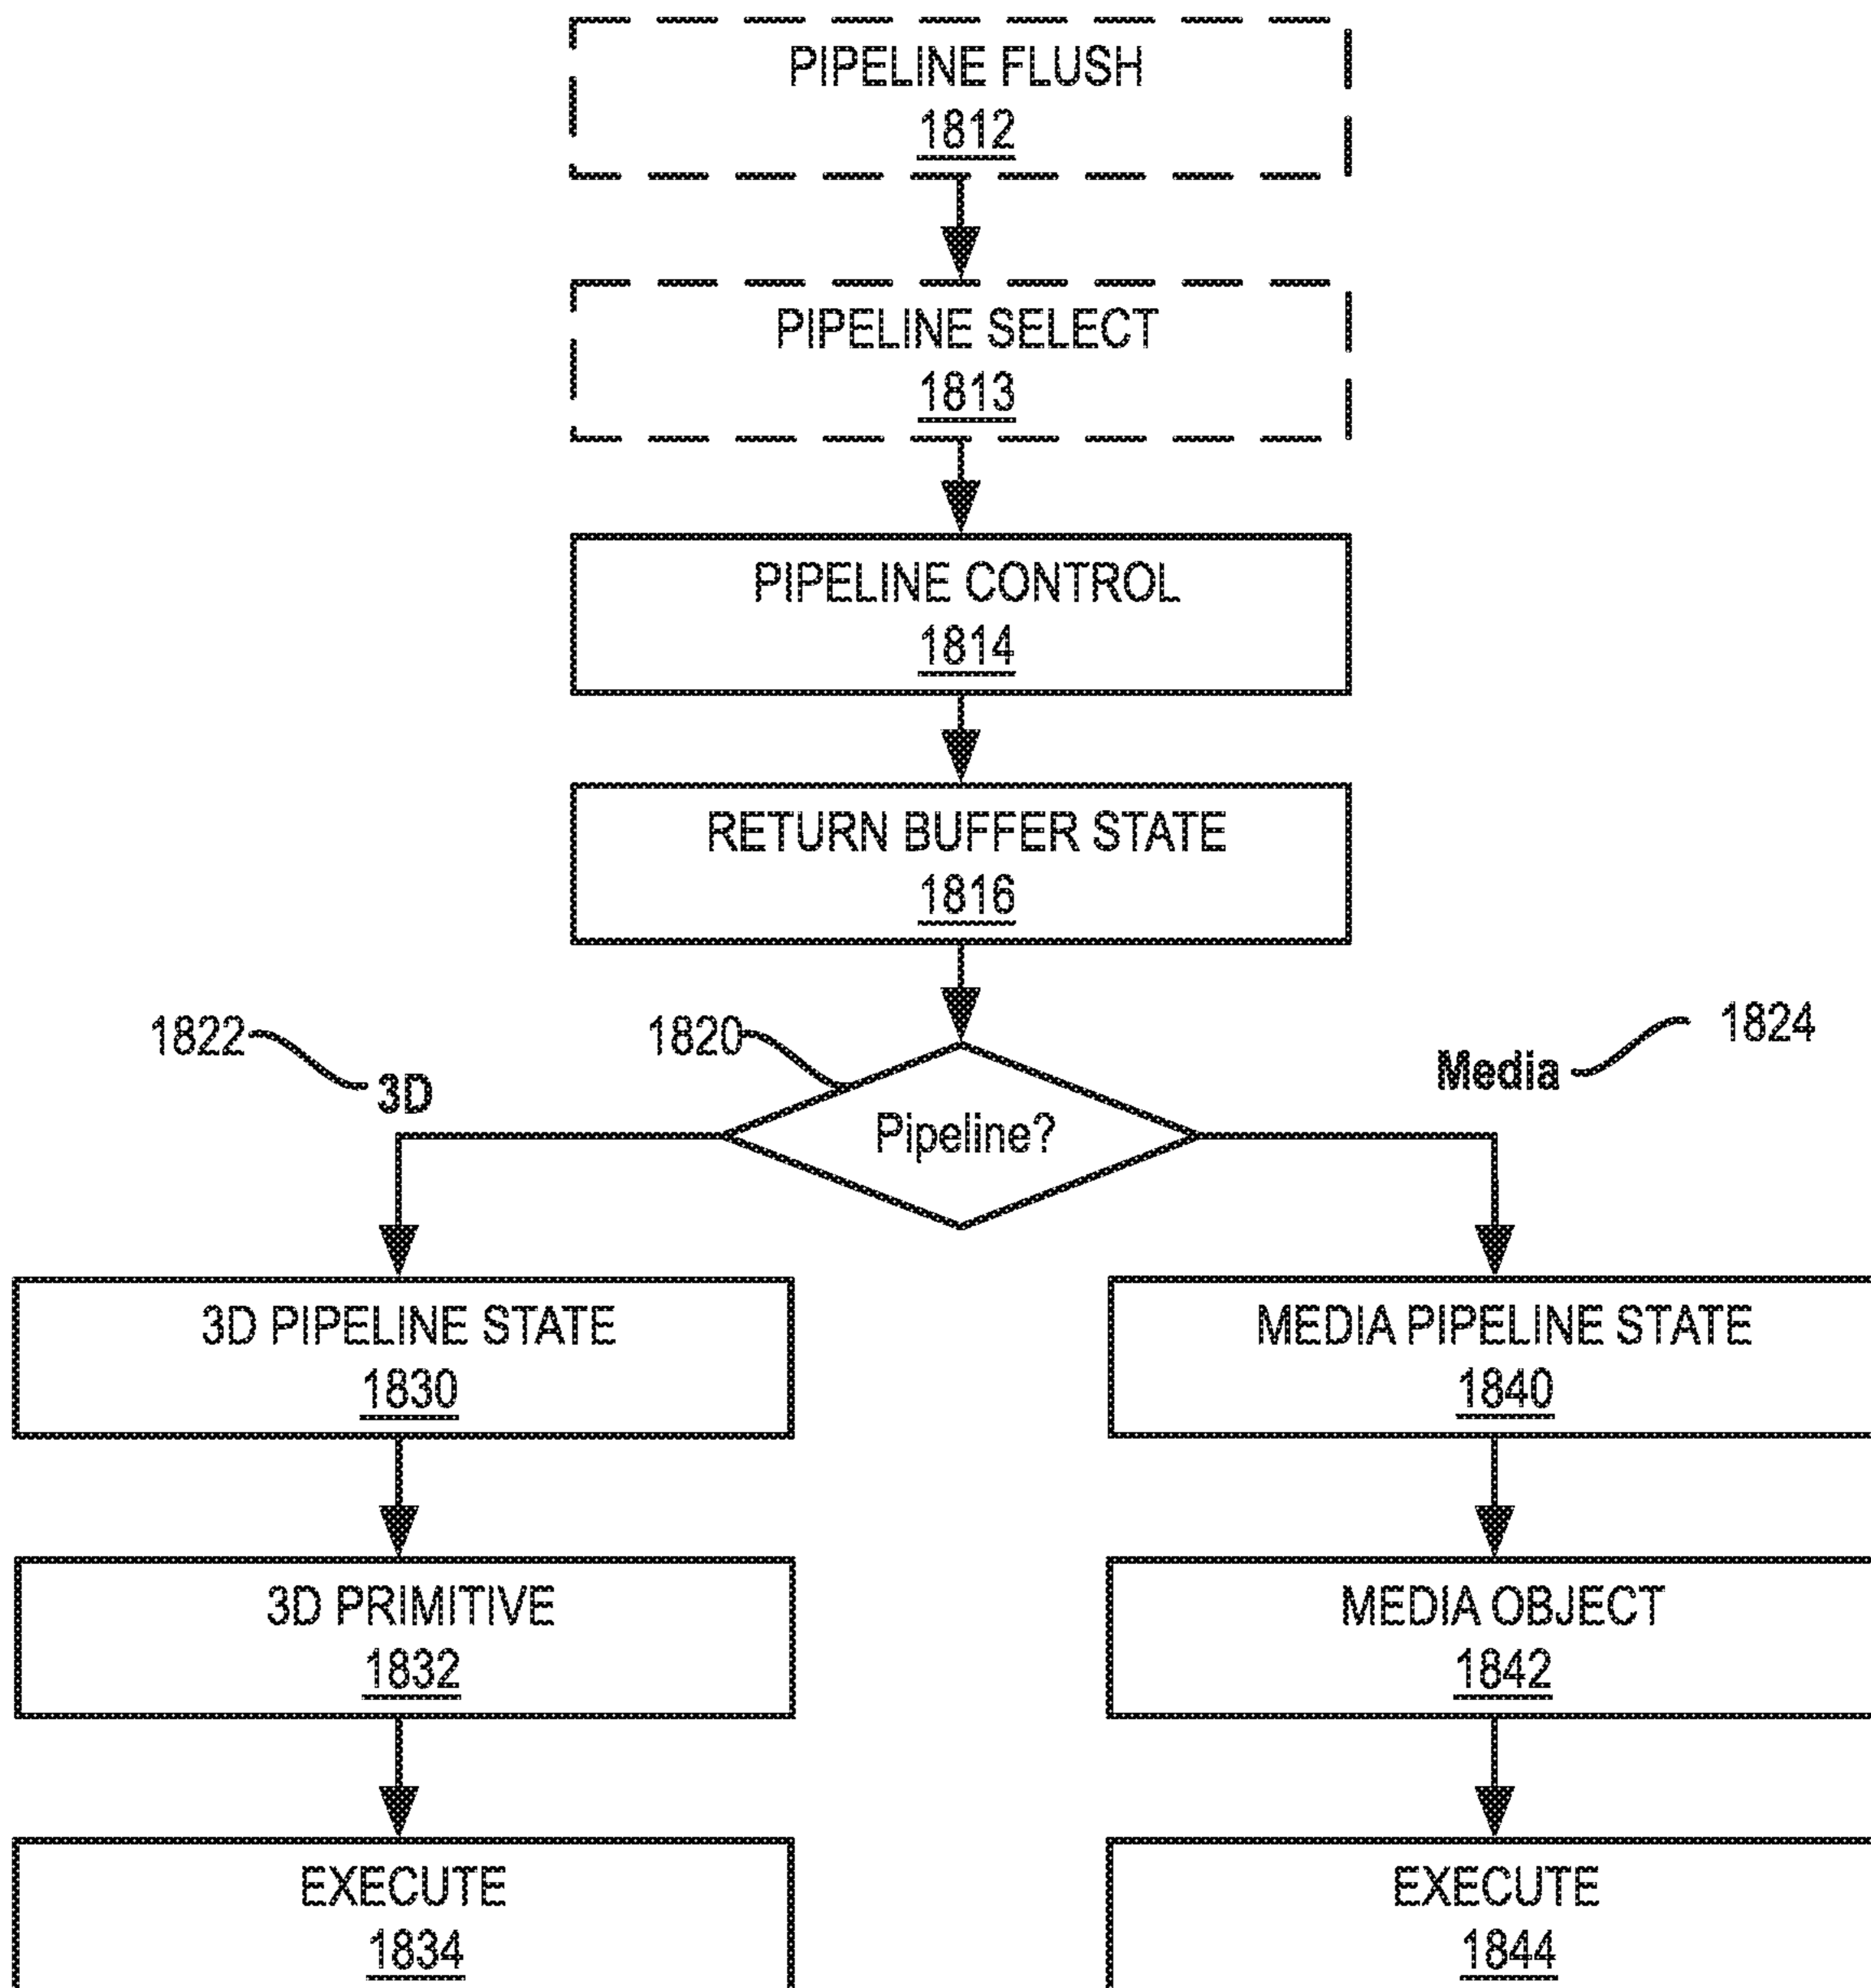


FIG. 18B

GRAPHICS PROCESSOR COMMAND SEQUENCE

1810



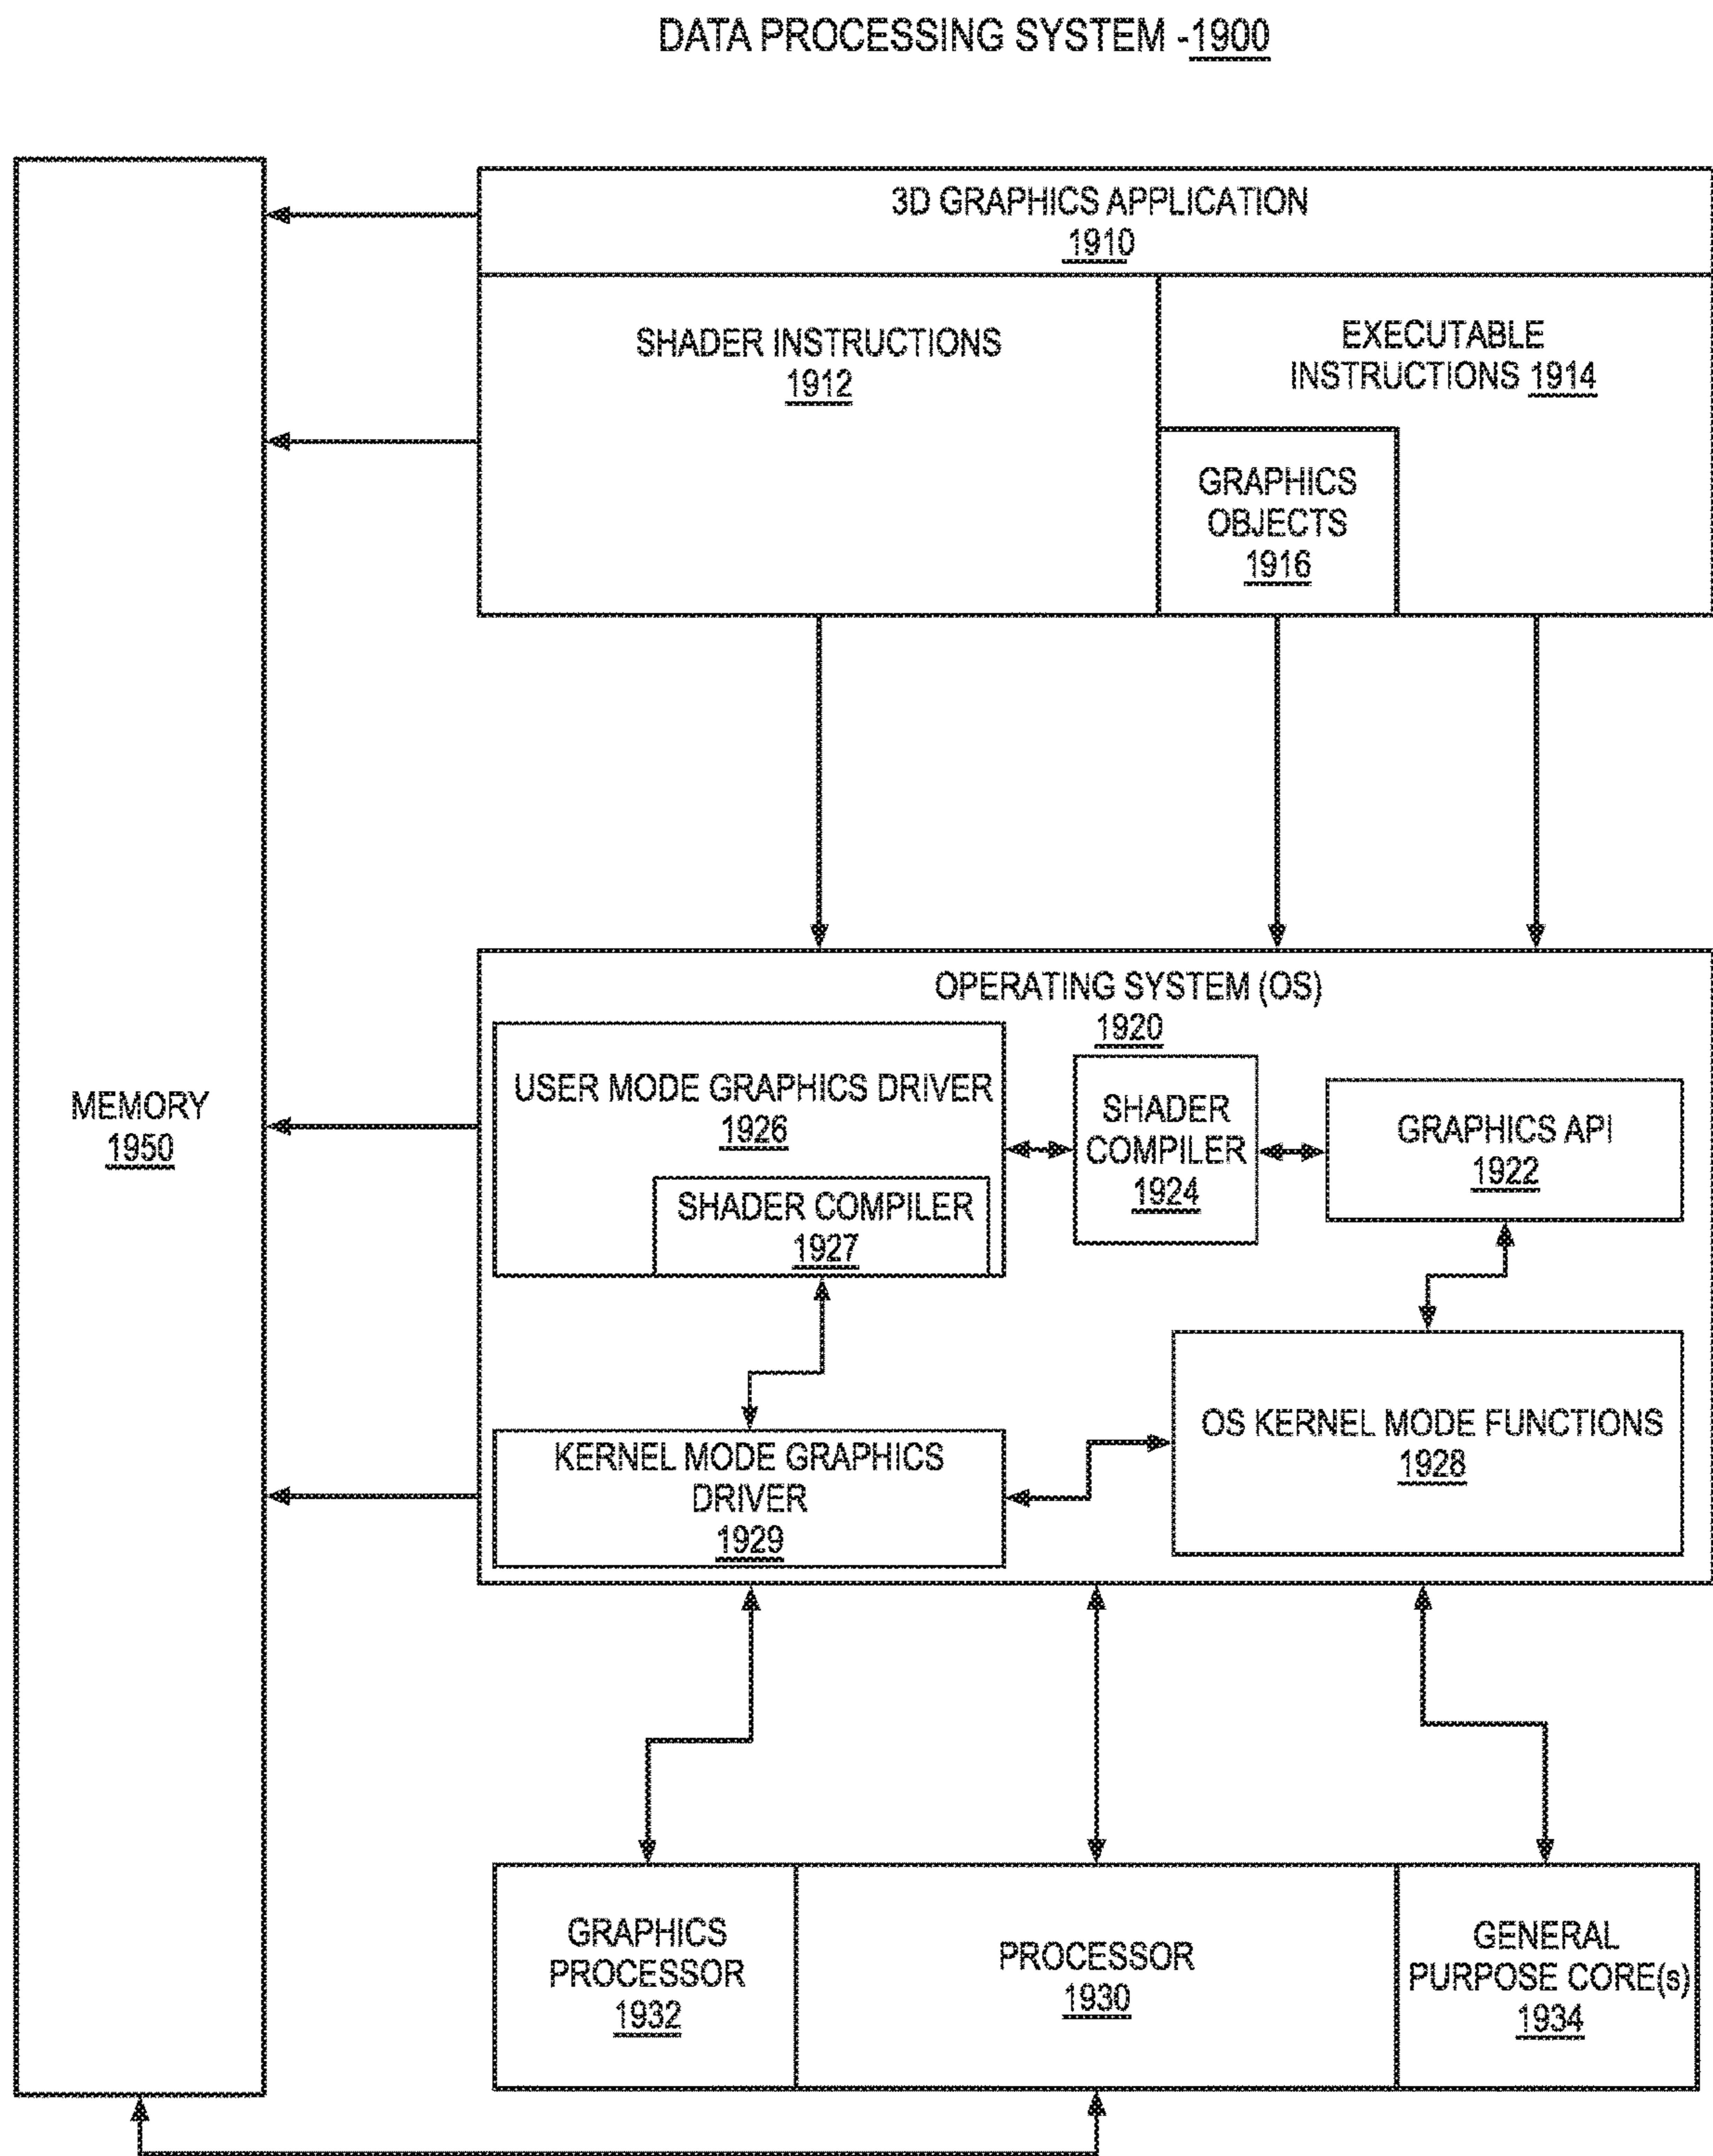


FIG. 19

IP CORE DEVELOPMENT - 2000




FIG. 20




FIG. 21




FIG. 22




FIG. 23

EFFICIENT MERGING OF ATOMIC OPERATIONS AT COMPUTING DEVICES

CLAIM TO PRIORITY

This Application is a continuation and claims priority of U.S. application Ser. No. 16/414,251, entitled EFFICIENT MERGING OF ATOMIC OPERATIONS AT COMPUTING DEVICES, by Joydeep Ray, et al., filed May 16, 2019 and issued as U.S. Pat. No. 10,796,401 on Oct. 6, 2020, which is a continuation and claims priority of U.S. application Ser. No. 15/488,565, entitled EFFICIENT MERGING OF ATOMIC OPERATIONS AT COMPUTING DEVICES, by Joydeep Ray, et al., filed Apr. 17, 2017 and issued as U.S. Pat. No. 10,325,344 on Jun. 18, 2019, the entire contents of which are incorporated herein by reference.

FIELD

Embodiments described herein relate generally to data processing and more particularly to facilitate efficient merging of atomic operations at computing devices.

BACKGROUND

Current parallel graphics data processing includes systems and methods developed to perform specific operations on graphics data such as, for example, linear interpolation, tessellation, rasterization, texture mapping, depth testing, etc. Traditionally, graphics processors used fixed function computational units to process graphics data; however, more recently, portions of graphics processors have been made programmable, enabling such processors to support a wider variety of operations for processing vertex and fragment data.

To further increase performance, graphics processors typically implement processing techniques such as pipelining that attempt to process, in parallel, as much graphics data as possible throughout the different parts of the graphics pipeline. Parallel graphics processors with single instruction, multiple thread (SIMT) architectures are designed to maximize the amount of parallel processing in the graphics pipeline. In an SIMT architecture, groups of parallel threads attempt to execute program instructions synchronously together as often as possible to increase processing efficiency. A general overview of software and hardware for SIMT architectures can be found in Shane Cook, *CUDA Programming*, Chapter 3, pages 37-51 (2013) and/or Nicholas Wilt, *CUDA Handbook, A Comprehensive Guide to GPU Programming*, Sections 2.6.2 to 3.1.2 (June 2013).

Conventional techniques process atomic operations serially (such as one at a time) when an address match is detected among single instruction, multiple data (SIMD) message slots. In existing graphics data port (GDP) or shared local memory (SLM) controller design, when multiple SIMD slots of an atomic message are mapped to the same address, atomic operations are serialized (such as processed one per cycle), which, in turn, results in an SIMD 16 atomic message (of 16 slots) can take up to 16 cycles to complete all the atomic operations. Such conventional techniques are inefficient and costly in terms of processing speed and system resources.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments are illustrated by way of example, and not by way of limitation, in the figures of the accompanying

drawings in which like reference numerals refer to similar elements. So that the manner in which the above recited features can be understood in detail, a more particular description, briefly summarized above, may be had by reference to embodiments, some of which are illustrated in the appended drawings. It is to be noted, however, that the appended drawings illustrate only typical embodiments and are therefore not to be considered limiting of its scope, for the drawings may illustrate other equally effective embodiments.

FIG. 1 is a block diagram illustrating a computing system configured to implement one or more embodiments.

FIG. 2A illustrates a parallel processor, according to an embodiment.

FIG. 2B is a block diagram of a partition unit, according to an embodiment.

FIG. 2C is a block diagram of a processing cluster 214 within a parallel processing unit, according to an embodiment.

FIG. 2D shows a graphics multiprocessor 234, according to one embodiment.

FIGS. 3A-3B illustrate additional graphics multiprocessors, according to embodiments.

FIG. 4A illustrates an exemplary architecture in which a plurality of GPU's are communicatively coupled to a plurality of multi-core processors 405-406 over high-speed links, according to an embodiment.

FIG. 4B illustrates additional details for an interconnection between a multicore processor and a graphics acceleration module, according to an embodiment.

FIG. 4C illustrates another embodiment in which the accelerator integration circuit 436 is integrated within the processor 407.

FIG. 4D illustrates an exemplary accelerator integration slice.

FIG. 4E illustrates additional details for one embodiment of a shared model.

FIG. 4F illustrates one embodiment employing a unified memory addressable via a common virtual memory address space used to access the physical processor memories and GPU memories.

FIG. 5 is a conceptual diagram of a graphics processing pipeline, according to an embodiment.

FIG. 6 illustrates a computing device hosting a dynamic atomic merging mechanism according to one embodiment.

FIG. 7 illustrates a dynamic atomic merging mechanism according to one embodiment.

FIG. 8 illustrates a conventional atomic pipeline.

FIG. 9A illustrates a novel atomic message pipeline allowing for dynamic merging of atomic operations according to one embodiment.

FIG. 9B illustrates an array of slots according to one embodiment.

FIG. 9C illustrates a method for merging of atomic operations for memory according to one embodiment.

FIG. 10 is a block diagram of an embodiment of a computer system with a processor having one or more processor cores and graphics processors.

FIG. 11 is a block diagram of one embodiment of a processor having one or more processor cores, an integrated memory controller, and an integrated graphics processor.

FIG. 12 is a block diagram of one embodiment of a graphics processor which may be a discreet graphics processing unit, or may be graphics processor integrated with a plurality of processing cores.

FIG. 13 is a block diagram of an embodiment of a graphics processing engine for a graphics processor.

FIG. 14 is a block diagram of another embodiment of a graphics processor.

FIG. 15 is a block diagram of thread execution logic including an array of processing elements.

FIG. 16 illustrates a graphics processor execution unit instruction format according to an embodiment.

FIG. 17 is a block diagram of another embodiment of a graphics processor which includes a graphics pipeline, a media pipeline, a display engine, thread execution logic, and a render output pipeline.

FIG. 18A is a block diagram illustrating a graphics processor command format according to an embodiment.

FIG. 18B is a block diagram illustrating a graphics processor command sequence according to an embodiment.

FIG. 19 illustrates exemplary graphics software architecture for a data processing system according to an embodiment.

FIG. 20 is a block diagram illustrating an IP core development system that may be used to manufacture an integrated circuit to perform operations according to an embodiment.

FIG. 21 is a block diagram illustrating an exemplary system on a chip integrated circuit that may be fabricated using one or more IP cores, according to an embodiment.

FIG. 22 is a block diagram illustrating an exemplary graphics processor of a system on a chip integrated circuit.

FIG. 23 is a block diagram illustrating an additional exemplary graphics processor of a system on a chip integrated circuit.

DETAILED DESCRIPTION

Embodiments provide for a novel technique to drastically speed up messages by merging individual atomic operations into a single. In one embodiment, when multiple slot addresses of an SIMD atomic message is the same, it is then detected based on an atomic operation and merges those atomic operations into a single atomic operation. For example, if an atomic operation is atomic increment (e.g., Atomic_increment) and, for example, all 16 slots addresses (of an SIMD 16) maps to the same address in memory, the final result may have the memory value incremented by 16. This novel technique provides for a mechanism to detect such scenarios and replace all the 16 atomic increments by a single atomic add operation (e.g., Atomic_ADD) by a value of 16.

In some embodiments, a graphics processing unit (GPU) is communicatively coupled to host/processor cores to accelerate graphics operations, machine-learning operations, pattern analysis operations, and various general purpose GPU (GPGPU) functions. The GPU may be communicatively coupled to the host processor/cores over a bus or another interconnect (e.g., a high-speed interconnect such as PCIe or NVLink). In other embodiments, the GPU may be integrated on the same package or chip as the cores and communicatively coupled to the cores over an internal processor bus/interconnect (i.e., internal to the package or chip). Regardless of the manner in which the GPU is connected, the processor cores may allocate work to the GPU in the form of sequences of commands/instructions contained in a work descriptor. The GPU then uses dedicated circuitry/logic for efficiently processing these commands/instructions.

In the following description, numerous specific details are set forth. However, embodiments, as described herein, may be practiced without these specific details. In other instances, well-known circuits, structures and techniques

have not been shown in details in order not to obscure the understanding of this description.

System Overview I

FIG. 1 is a block diagram illustrating a computing system 100 configured to implement one or more aspects of the embodiments described herein. The computing system 100 includes a processing subsystem 101 having one or more processor(s) 102 and a system memory 104 communicating via an interconnection path that may include a memory hub 105. The memory hub 105 may be a separate component within a chipset component or may be integrated within the one or more processor(s) 102. The memory hub 105 couples with an I/O subsystem 111 via a communication link 106. The I/O subsystem 111 includes an I/O hub 107 that can enable the computing system 100 to receive input from one or more input device(s) 108. Additionally, the I/O hub 107 can enable a display controller, which may be included in the one or more processor(s) 102, to provide outputs to one or more display device(s) 110A. In one embodiment, the one or more display device(s) 110A coupled with the I/O hub 107 can include a local, internal, or embedded display device.

In one embodiment, the processing subsystem 101 includes one or more parallel processor(s) 112 coupled to memory hub 105 via a bus or other communication link 113. The communication link 113 may be one of any number of standards based communication link technologies or protocols, such as, but not limited to PCI Express, or may be a vendor specific communications interface or communications fabric. In one embodiment, the one or more parallel processor(s) 112 form a computationally focused parallel or vector processing system that include a large number of processing cores and/or processing clusters, such as a many integrated core (MIC) processor. In one embodiment, the one or more parallel processor(s) 112 form a graphics processing subsystem that can output pixels to one of the one or more display device(s) 110A coupled via the I/O Hub 107. The one or more parallel processor(s) 112 can also include a display controller and display interface (not shown) to enable a direct connection to one or more display device(s) 110B.

Within the I/O subsystem 111, a system storage unit 114 can connect to the I/O hub 107 to provide a storage mechanism for the computing system 100. An I/O switch 116 can be used to provide an interface mechanism to enable connections between the I/O hub 107 and other components, such as a network adapter 118 and/or wireless network adapter 119 that may be integrated into the platform, and various other devices that can be added via one or more add-in device(s) 120. The network adapter 118 can be an Ethernet adapter or another wired network adapter. The wireless network adapter 119 can include one or more of a Wi-Fi, Bluetooth, near field communication (NFC), or other network device that includes one or more wireless radios.

The computing system 100 can include other components not explicitly shown, including USB or other port connections, optical storage drives, video capture devices, and the like, may also be connected to the I/O hub 107. Communication paths interconnecting the various components in FIG. 1 may be implemented using any suitable protocols, such as PCI (Peripheral Component Interconnect) based protocols (e.g., PCI-Express), or any other bus or point-to-point communication interfaces and/or protocol(s), such as the NV-Link high-speed interconnect, or interconnect protocols known in the art.

In one embodiment, the one or more parallel processor(s) 112 incorporate circuitry optimized for graphics and video processing, including, for example, video output circuitry,

and constitutes a graphics processing unit (GPU). In another embodiment, the one or more parallel processor(s) 112 incorporate circuitry optimized for general purpose processing, while preserving the underlying computational architecture, described in greater detail herein. In yet another embodiment, components of the computing system 100 may be integrated with one or more other system elements on a single integrated circuit. For example, the one or more parallel processor(s), 112 memory hub 105, processor(s) 102, and I/O hub 107 can be integrated into a system on chip (SoC) integrated circuit. Alternatively, the components of the computing system 100 can be integrated into a single package to form a system in package (SIP) configuration. In one embodiment, at least a portion of the components of the computing system 100 can be integrated into a multi-chip module (MCM), which can be interconnected with other multi-chip modules into a modular computing system.

It will be appreciated that the computing system 100 shown herein is illustrative and that variations and modifications are possible. The connection topology, including the number and arrangement of bridges, the number of processor(s) 102, and the number of parallel processor(s) 112, may be modified as desired. For instance, in some embodiments, system memory 104 is connected to the processor(s) 102 directly rather than through a bridge, while other devices communicate with system memory 104 via the memory hub 105 and the processor(s) 102. In other alternative topologies, the parallel processor(s) 112 are connected to the I/O hub 107 or directly to one of the one or more processor(s) 102, rather than to the memory hub 105. In other embodiments, the I/O hub 107 and memory hub 105 may be integrated into a single chip. Some embodiments may include two or more sets of processor(s) 102 attached via multiple sockets, which can couple with two or more instances of the parallel processor(s) 112.

Some of the particular components shown herein are optional and may not be included in all implementations of the computing system 100. For example, any number of add-in cards or peripherals may be supported, or some components may be eliminated. Furthermore, some architectures may use different terminology for components similar to those illustrated in FIG. 1. For example, the memory hub 105 may be referred to as a Northbridge in some architectures, while the I/O hub 107 may be referred to as a Southbridge.

FIG. 2A illustrates a parallel processor 200, according to an embodiment. The various components of the parallel processor 200 may be implemented using one or more integrated circuit devices, such as programmable processors, application specific integrated circuits (ASICs), or field programmable gate arrays (FPGA). The illustrated parallel processor 200 is a variant of the one or more parallel processor(s) 112 shown in FIG. 1, according to an embodiment.

In one embodiment, the parallel processor 200 includes a parallel processing unit 202. The parallel processing unit includes an I/O unit 204 that enables communication with other devices, including other instances of the parallel processing unit 202. The I/O unit 204 may be directly connected to other devices. In one embodiment, the I/O unit 204 connects with other devices via the use of a hub or switch interface, such as memory hub 105. The connections between the memory hub 105 and the I/O unit 204 form a communication link 113. Within the parallel processing unit 202, the I/O unit 204 connects with a host interface 206 and a memory crossbar 216, where the host interface 206 receives commands directed to performing processing

operations and the memory crossbar 216 receives commands directed to performing memory operations.

When the host interface 206 receives a command buffer via the I/O unit 204, the host interface 206 can direct work operations to perform those commands to a front end 208. In one embodiment, the front end 208 couples with a scheduler 210, which is configured to distribute commands or other work items to a processing cluster array 212. In one embodiment, the scheduler 210 ensures that the processing cluster array 212 is properly configured and in a valid state before tasks are distributed to the processing clusters of the processing cluster array 212.

The processing cluster array 212 can include up to "N" processing clusters (e.g., cluster 214A, cluster 214B, through cluster 214N). Each cluster 214A-214N of the processing cluster array 212 can execute a large number of concurrent threads. The scheduler 210 can allocate work to the clusters 214A-214N of the processing cluster array 212 using various scheduling and/or work distribution algorithms, which may vary depending on the workload arising for each type of program or computation. The scheduling can be handled dynamically by the scheduler 210, or can be assisted in part by compiler logic during compilation of program logic configured for execution by the processing cluster array 212.

In one embodiment, different clusters 214A-214N of processing cluster array 212 can be allocated for processing different types of programs or for performing different types of computations.

The processing cluster array 212 can be configured to perform various types of parallel processing operations. In one embodiment, the processing cluster array 212 is configured to perform general-purpose parallel compute operations. For example, the processing cluster array 212 can include logic to execute processing tasks including filtering of video and/or audio data, performing modeling operations, including physics operations, and performing data transformations.

In one embodiment, the processing cluster array 212 is configured to perform parallel graphics processing operations. In embodiments in which the parallel processor 200 is configured to perform graphics processing operations, the processing cluster array 212 can include additional logic to support the execution of such graphics processing operations, including, but not limited to texture sampling logic to perform texture operations, as well as tessellation logic and other vertex processing logic. Additionally, the processing cluster array 212 can be configured to execute graphics processing related shader programs such as, but not limited to vertex shaders, tessellation shaders, geometry shaders, and pixel shaders. The parallel processing unit 202 can transfer data from system memory via the I/O unit 204 for processing. During processing the transferred data can be stored to on-chip memory (e.g., parallel processor memory 222) during processing, then written back to system memory.

In one embodiment, when the parallel processing unit 202 is used to perform graphics processing, the scheduler 210 can be configured to divide the processing workload into approximately equal sized tasks, to better enable distribution of the graphics processing operations to multiple clusters 214A-214N of the processing cluster array 212. In some embodiments, portions of the processing cluster array 212 can be configured to perform different types of processing. For example, a first portion may be configured to perform vertex shading and topology generation, a second portion may be configured to perform tessellation and geometry

shading, and a third portion may be configured to perform pixel shading or other screen space operations, to produce a rendered image for display. Intermediate data produced by one or more of the clusters **214A-214N** may be stored in buffers to allow the intermediate data to be transmitted between clusters **214A-214N** for further processing.

During operation, the processing cluster array **212** can receive processing tasks to be executed via the scheduler **210**, which receives commands defining processing tasks from front end **208**. For graphics processing operations, processing tasks can include indices of data to be processed, e.g., surface (patch) data, primitive data, vertex data, and/or pixel data, as well as state parameters and commands defining how the data is to be processed (e.g., what program is to be executed). The scheduler **210** may be configured to fetch the indices corresponding to the tasks or may receive the indices from the front end **208**. The front end **208** can be configured to ensure the processing cluster array **212** is configured to a valid state before the workload specified by incoming command buffers (e.g., batch-buffers, push buffers, etc.) is initiated.

Each of the one or more instances of the parallel processing unit **202** can couple with parallel processor memory **222**. The parallel processor memory **222** can be accessed via the memory crossbar **216**, which can receive memory requests from the processing cluster array **212** as well as the I/O unit **204**. The memory crossbar **216** can access the parallel processor memory **222** via a memory interface **218**. The memory interface **218** can include multiple partition units (e.g., partition unit **220A**, partition unit **220B**, through partition unit **220N**) that can each couple to a portion (e.g., memory unit) of parallel processor memory **222**. In one implementation, the number of partition units **220A-220N** is configured to be equal to the number of memory units, such that a first partition unit **220A** has a corresponding first memory unit **224A**, a second partition unit **220B** has a corresponding memory unit **224B**, and an Nth partition unit **220N** has a corresponding Nth memory unit **224N**. In other embodiments, the number of partition units **220A-220N** may not be equal to the number of memory devices.

In various embodiments, the memory units **224A-224N** can include various types of memory devices, including dynamic random access memory (DRAM) or graphics random access memory, such as synchronous graphics random access memory (SGRAM), including graphics double data rate (GDDR) memory. In one embodiment, the memory units **224A-224N** may also include 3D stacked memory, including but not limited to high bandwidth memory (HBM). Persons skilled in the art will appreciate that the specific implementation of the memory units **224A-224N** can vary, and can be selected from one of various conventional designs. Render targets, such as frame buffers or texture maps may be stored across the memory units **224A-224N**, allowing partition units **220A-220N** to write portions of each render target in parallel to efficiently use the available bandwidth of parallel processor memory **222**. In some embodiments, a local instance of the parallel processor memory **222** may be excluded in favor of a unified memory design that utilizes system memory in conjunction with local cache memory.

In one embodiment, any one of the clusters **214A-214N** of the processing cluster array **212** can process data that will be written to any of the memory units **224A-224N** within parallel processor memory **222**. The memory crossbar **216** can be configured to transfer the output of each cluster **214A-214N** to any partition unit **220A-220N** or to another cluster **214A-214N**, which can perform additional process-

ing operations on the output. Each cluster **214A-214N** can communicate with the memory interface **218** through the memory crossbar **216** to read from or write to various external memory devices. In one embodiment, the memory crossbar **216** has a connection to the memory interface **218** to communicate with the I/O unit **204**, as well as a connection to a local instance of the parallel processor memory **222**, enabling the processing units within the different processing clusters **214A-214N** to communicate with system memory or other memory that is not local to the parallel processing unit **202**. In one embodiment, the memory crossbar **216** can use virtual channels to separate traffic streams between the clusters **214A-214N** and the partition units **220A-220N**.

While a single instance of the parallel processing unit **202** is illustrated within the parallel processor **200**, any number of instances of the parallel processing unit **202** can be included. For example, multiple instances of the parallel processing unit **202** can be provided on a single add-in card, or multiple add-in cards can be interconnected. The different instances of the parallel processing unit **202** can be configured to inter-operate even if the different instances have different numbers of processing cores, different amounts of local parallel processor memory, and/or other configuration differences. For example, and in one embodiment, some instances of the parallel processing unit **202** can include higher precision floating point units relative to other instances. Systems incorporating one or more instances of the parallel processing unit **202** or the parallel processor **200** can be implemented in a variety of configurations and form factors, including but not limited to desktop, laptop, or handheld personal computers, servers, workstations, game consoles, and/or embedded systems.

FIG. **2B** is a block diagram of a partition unit **220**, according to an embodiment. In one embodiment, the partition unit **220** is an instance of one of the partition units **220A-220N** of FIG. **2A**. As illustrated, the partition unit **220** includes an L2 cache **221**, a frame buffer interface **225**, and a ROP **226** (raster operations unit). The L2 cache **221** is a read/write cache that is configured to perform load and store operations received from the memory crossbar **216** and ROP **226**. Read misses and urgent write-back requests are output by L2 cache **221** to frame buffer interface **225** for processing. Dirty updates can also be sent to the frame buffer via the frame buffer interface **225** for opportunistic processing. In one embodiment, the frame buffer interface **225** interfaces with one of the memory units in parallel processor memory, such as the memory units **224A-224N** of FIG. **2A** (e.g., within parallel processor memory **222**).

In graphics applications, the ROP **226** is a processing unit that performs raster operations, such as stencil, z test, blending, and the like. The ROP **226** then outputs processed graphics data that is stored in graphics memory. In some embodiments, the ROP **226** includes compression logic to compress z or color data that is written to memory and decompress z or color data that is read from memory. In some embodiments, the ROP **226** is included within each processing cluster (e.g., cluster **214A-214N** of FIG. **2A**) instead of within the partition unit **220**. In such embodiment, read and write requests for pixel data are transmitted over the memory crossbar **216** instead of pixel fragment data.

The processed graphics data may be displayed on a display device, such as one of the one or more display device(s) **110** of FIG. **1**, routed for further processing by the processor(s) **102**, or routed for further processing by one of the processing entities within the parallel processor **200** of FIG. **2A**.

FIG. 2C is a block diagram of a processing cluster **214** within a parallel processing unit, according to an embodiment. In one embodiment, the processing cluster is an instance of one of the processing clusters **214A-214N** of FIG. 2A. The processing cluster **214** can be configured to execute many threads in parallel, where the term “thread” refers to an instance of a particular program executing on a particular set of input data. In some embodiments, single-instruction, multiple-data (SIMD) instruction issue techniques are used to support parallel execution of a large number of threads without providing multiple independent instruction units. In other embodiments, single-instruction, multiple-thread (SIMT) techniques are used to support parallel execution of a large number of generally synchronized threads, using a common instruction unit configured to issue instructions to a set of processing engines within each one of the processing clusters. Unlike a SIMD execution regime, where all processing engines typically execute identical instructions, SIMT execution allows different threads to more readily follow divergent execution paths through a given thread program. Persons skilled in the art will understand that a SIMD processing regime represents a functional subset of a SIMT processing regime.

Operation of the processing cluster **214** can be controlled via a pipeline manager **232** that distributes processing tasks to SIMT parallel processors. The pipeline manager **232** receives instructions from the scheduler **210** of FIG. 2A and manages execution of those instructions via a graphics multiprocessor **234** and/or a texture unit **236**. The illustrated graphics multiprocessor **234** is an exemplary instance of an SIMT parallel processor. However, various types of SIMT parallel processors of differing architectures may be included within the processing cluster **214**. One or more instances of the graphics multiprocessor **234** can be included within a processing cluster **214**. The graphics multiprocessor **234** can process data and a data crossbar **240** can be used to distribute the processed data to one of multiple possible destinations, including other shader units. The pipeline manager **232** can facilitate the distribution of processed data by specifying destinations for processed data to be distributed via the data crossbar **240**.

Each graphics multiprocessor **234** within the processing cluster **214** can include an identical set of functional execution logic (e.g., arithmetic logic units, load-store units, etc.). The functional execution logic can be configured in a pipelined manner in which new instructions can be issued before previous instructions are complete. The functional execution logic may be provided. The functional logic supports a variety of operations including integer and floating point arithmetic comparison operations, Boolean operations bit-shifting, and computation of various algebraic functions. In one embodiment, the same functional-unit hardware can be leveraged to perform different operations and any combination of functional units may be present.

The instructions transmitted to the processing cluster **214** constitutes a thread. A set of threads executing across the set of parallel processing engines is a thread group. A thread group executes the same program on different input data. Each thread within a thread group can be assigned to a different processing engine within a graphics multiprocessor **234**. A thread group may include fewer threads than the number of processing engines within the graphics multiprocessor **234**. When a thread group includes fewer threads than the number of processing engines, one or more of the processing engines may be idle during cycles in which that thread group is being processed. A thread group may also include more threads than the number of processing engines

within the graphics multiprocessor **234**. When the thread group includes more threads than the number of processing engines within the graphics multiprocessor **234**, processing can be performed over consecutive clock cycles. In one embodiment, multiple thread groups can be executed concurrently on a graphics multiprocessor **234**.

In one embodiment, the graphics multiprocessor **234** includes an internal cache memory to perform load and store operations. In one embodiment, the graphics multiprocessor **234** can forego an internal cache and use a cache memory (e.g., L1 cache **308**) within the processing cluster **214**. Each graphics multiprocessor **234** also has access to L2 caches within the partition units (e.g., partition units **220A-220N** of FIG. 2A) that are shared among all processing clusters **214** and may be used to transfer data between threads. The graphics multiprocessor **234** may also access off-chip global memory, which can include one or more of local parallel processor memory and/or system memory. Any memory external to the parallel processing unit **202** may be used as global memory. Embodiments in which the processing cluster **214** includes multiple instances of the graphics multiprocessor **234** can share common instructions and data, which may be stored in the L1 cache **308**.

Each processing cluster **214** may include an MMU **245** (memory management unit) that is configured to map virtual addresses into physical addresses. In other embodiments, one or more instances of the MMU **245** may reside within the memory interface **218** of FIG. 2A. The MMU **245** includes a set of page table entries (PTEs) used to map a virtual address to a physical address of a tile (talk more about tiling) and optionally a cache line index. The MMU **245** may include address translation lookaside buffers (TLB) or caches that may reside within the graphics multiprocessor **234** or the L1 cache or processing cluster **214**. The physical address is processed to distribute surface data access locality to allow efficient request interleaving among partition units. The cache line index may be used to determine whether a request for a cache line is a hit or miss.

In graphics and computing applications, a processing cluster **214** may be configured such that each graphics multiprocessor **234** is coupled to a texture unit **236** for performing texture mapping operations, e.g., determining texture sample positions, reading texture data, and filtering the texture data. Texture data is read from an internal texture L1 cache (not shown) or in some embodiments from the L1 cache within graphics multiprocessor **234** and is fetched from an L2 cache, local parallel processor memory, or system memory, as needed. Each graphics multiprocessor **234** outputs processed tasks to the data crossbar **240** to provide the processed task to another processing cluster **214** for further processing or to store the processed task in an L2 cache, local parallel processor memory, or system memory via the memory crossbar **216**. A preROP **242** (pre-raster operations unit) is configured to receive data from graphics multiprocessor **234**, direct data to ROP units, which may be located with partition units as described herein (e.g., partition units **220A-220N** of FIG. 2A). The preROP **242** unit can perform optimizations for color blending, organize pixel color data, and perform address translations.

It will be appreciated that the core architecture described herein is illustrative and that variations and modifications are possible. Any number of processing units, e.g., graphics multiprocessor **234**, texture units **236**, preROPs **242**, etc., may be included within a processing cluster **214**. Further, while only one processing cluster **214** is shown, a parallel processing unit as described herein may include any number of instances of the processing cluster **214**. In one embodi-

ment, each processing cluster **214** can be configured to operate independently of other processing clusters **214** using separate and distinct processing units, L1 caches, etc.

FIG. **2D** shows a graphics multiprocessor **234**, according to one embodiment. In such embodiment, the graphics multiprocessor **234** couples with the pipeline manager **232** of the processing cluster **214**. The graphics multiprocessor **234** has an execution pipeline including but not limited to an instruction cache **252**, an instruction unit **254**, an address mapping unit **256**, a register file **258**, one or more general purpose graphics processing unit (GPGPU) cores **262**, and one or more load/store units **266**. The GPGPU cores **262** and load/store units **266** are coupled with cache memory **272** and shared memory **270** via a memory and cache interconnect **268**.

In one embodiment, the instruction cache **252** receives a stream of instructions to execute from the pipeline manager **232**. The instructions are cached in the instruction cache **252** and dispatched for execution by the instruction unit **254**. The instruction unit **254** can dispatch instructions as thread groups (e.g., warps), with each thread of the thread group assigned to a different execution unit within GPGPU core **262**. An instruction can access any of a local, shared, or global address space by specifying an address within a unified address space. The address mapping unit **256** can be used to translate addresses in the unified address space into a distinct memory address that can be accessed by the load/store units **266**.

The register file **258** provides a set of registers for the functional units of the graphics multiprocessor **324**. The register file **258** provides temporary storage for operands connected to the data paths of the functional units (e.g., GPGPU cores **262**, load/store units **266**) of the graphics multiprocessor **324**. In one embodiment, the register file **258** is divided between each of the functional units such that each functional unit is allocated a dedicated portion of the register file **258**. In one embodiment, the register file **258** is divided between the different warps being executed by the graphics multiprocessor **324**.

The GPGPU cores **262** can each include floating point units (FPUs) and/or integer arithmetic logic units (ALUs) that are used to execute instructions of the graphics multiprocessor **324**. The GPGPU cores **262** can be similar in architecture or can differ in architecture, according to embodiments. For example, and in one embodiment, a first portion of the GPGPU cores **262** include a single precision FPU and an integer ALU while a second portion of the GPGPU cores include a double precision FPU. In one embodiment, the FPUs can implement the IEEE 754-2008 standard for floating point arithmetic or enable variable precision floating point arithmetic. The graphics multiprocessor **324** can additionally include one or more fixed function or special function units to perform specific functions such as copy rectangle or pixel blending operations. In one embodiment one or more of the GPGPU cores can also include fixed or special function logic.

The memory and cache interconnect **268** is an interconnect network that connects each of the functional units of the graphics multiprocessor **324** to the register file **258** and to the shared memory **270**. In one embodiment, the memory and cache interconnect **268** is a crossbar interconnect that allows the load/store unit **266** to implement load and store operations between the shared memory **270** and the register file **258**. The register file **258** can operate at the same frequency as the GPGPU cores **262**, thus data transfer between the GPGPU cores **262** and the register file **258** is very low latency. The shared memory **270** can be used to

enable communication between threads that execute on the functional units within the graphics multiprocessor **234**. The cache memory **272** can be used as a data cache for example, to cache texture data communicated between the functional units and the texture unit **236**. The shared memory **270** can also be used as a program managed cache. Threads executing on the GPGPU cores **262** can programmatically store data within the shared memory in addition to the automatically cached data that is stored within the cache memory **272**.

FIGS. **3A-3B** illustrate additional graphics multiprocessors, according to embodiments. The illustrated graphics multiprocessors **325**, **350** are variants of the graphics multiprocessor **234** of FIG. **2C**. The illustrated graphics multiprocessors **325**, **350** can be configured as a streaming multiprocessor (SM) capable of simultaneous execution of a large number of execution threads.

FIG. **3A** shows a graphics multiprocessor **325** according to an additional embodiment. The graphics multiprocessor **325** includes multiple additional instances of execution resource units relative to the graphics multiprocessor **234** of FIG. **2D**. For example, the graphics multiprocessor **325** can include multiple instances of the instruction unit **332A-332B**, register file **334A-334B**, and texture unit(s) **344A-344B**. The graphics multiprocessor **325** also includes multiple sets of graphics or compute execution units (e.g., GPGPU core **336A-336B**, GPGPU core **337A-337B**, GPGPU core **338A-338B**) and multiple sets of load/store units **340A-340B**. In one embodiment, the execution resource units have a common instruction cache **330**, texture and/or data cache memory **342**, and shared memory **346**. The various components can communicate via an interconnect fabric **327**. In one embodiment, the interconnect fabric **327** includes one or more crossbar switches to enable communication between the various components of the graphics multiprocessor **325**.

FIG. **3B** shows a graphics multiprocessor **350** according to an additional embodiment. The graphics processor includes multiple sets of execution resources **356A-356D**, where each set of execution resource includes multiple instruction units, register files, GPGPU cores, and load store units, as illustrated in FIG. **2D** and FIG. **3A**. The execution resources **356A-356D** can work in concert with texture unit(s) **360A-360D** for texture operations, while sharing an instruction cache **354**, and shared memory **362**. In one embodiment, the execution resources **356A-356D** can share an instruction cache **354** and shared memory **362**, as well as multiple instances of a texture and/or data cache memory **358A-358B**. The various components can communicate via an interconnect fabric **352** similar to the interconnect fabric **327** of FIG. **3A**.

Persons skilled in the art will understand that the architecture described in FIGS. **1**, **2A-2D**, and **3A-3B** are descriptive and not limiting as to the scope of the present embodiments. Thus, the techniques described herein may be implemented on any properly configured processing unit, including, without limitation, one or more mobile application processors, one or more desktop or server central processing units (CPUs) including multi-core CPUs, one or more parallel processing units, such as the parallel processing unit **202** of FIG. **2A**, as well as one or more graphics processors or special purpose processing units, without departure from the scope of the embodiments described herein.

In some embodiments, a parallel processor or GPGPU as described herein is communicatively coupled to host/processor cores to accelerate graphics operations, machine-

learning operations, pattern analysis operations, and various general purpose GPU (GPGPU) functions. The GPU may be communicatively coupled to the host processor/cores over a bus or other interconnect (e.g., a high-speed interconnect such as PCIe or NVLink). In other embodiments, the GPU may be integrated on the same package or chip as the cores and communicatively coupled to the cores over an internal processor bus/interconnect (i.e., internal to the package or chip). Regardless of the manner in which the GPU is connected, the processor cores may allocate work to the GPU in the form of sequences of commands/instructions contained in a work descriptor. The GPU then uses dedicated circuitry/logic for efficiently processing these commands/instructions.

Techniques for GPU to Host Processor Interconnection

FIG. 4A illustrates an exemplary architecture in which a plurality of GPUs **410-413** are communicatively coupled to a plurality of multi-core processors **405-406** over high-speed links **440-443** (e.g., buses, point-to-point interconnects, etc.). In one embodiment, the high-speed links **440-443** support a communication throughput of 4 GB/s, 30 GB/s, 80 GB/s or higher, depending on the implementation. Various interconnect protocols may be used including, but not limited to, PCIe 4.0 or 5.0 and NVLink 2.0. However, the underlying principles of the invention are not limited to any particular communication protocol or throughput.

In addition, in one embodiment, two or more of the GPUs **410-413** are interconnected over high-speed links **444-445**, which may be implemented using the same or different protocols/links than those used for high-speed links **440-443**. Similarly, two or more of the multi-core processors **405-406** may be connected over high speed link **433** which may be symmetric multi-processor (SMP) buses operating at 20 GB/s, 30 GB/s, 120 GB/s or higher. Alternatively, all communication between the various system components shown in FIG. 4A may be accomplished using the same protocols/links (e.g., over a common interconnection fabric). As mentioned, however, the underlying principles of the invention are not limited to any particular type of interconnect technology.

In one embodiment, each multi-core processor **405-406** is communicatively coupled to a processor memory **401-402**, via memory interconnects **430-431**, respectively, and each GPU **410-413** is communicatively coupled to GPU memory **420-423** over GPU memory interconnects **450-453**, respectively. The memory interconnects **430-431** and **450-453** may utilize the same or different memory access technologies. By way of example, and not limitation, the processor memories **401-402** and GPU memories **420-423** may be volatile memories such as dynamic random access memories (DRAMs) (including stacked DRAMs), Graphics DDR SDRAM (GDDR) (e.g., GDDR5, GDDR6), or High Bandwidth Memory (HBM) and/or may be non-volatile memories such as 3D XPoint or Nano-Ram. In one embodiment, some portion of the memories may be volatile memory and another portion may be non-volatile memory (e.g., using a two-level memory (2LM) hierarchy).

As described below, although the various processors **405-406** and GPUs **410-413** may be physically coupled to a particular memory **401-402**, **420-423**, respectively, a unified memory architecture may be implemented in which the same virtual system address space (also referred to as the “effective address” space) is distributed among all of the various physical memories. For example, processor memories **401-402** may each comprise 64 GB of the system memory address space and GPU memories **420-423** may

each comprise 32 GB of the system memory address space (resulting in a total of 256 GB addressable memory in this example).

FIG. 4B illustrates additional details for an interconnection between a multi-core processor **407** and a graphics acceleration module **446** in accordance with one embodiment. The graphics acceleration module **446** may include one or more GPU chips integrated on a line card which is coupled to the processor **407** via the high-speed link **440**. Alternatively, the graphics acceleration module **446** may be integrated on the same package or chip as the processor **407**.

The illustrated processor **407** includes a plurality of cores **460A-460D**, each with a translation lookaside buffer **461A-461D** and one or more caches **462A-462D**. The cores may include various other components for executing instructions and processing data which are not illustrated to avoid obscuring the underlying principles of the invention (e.g., instruction fetch units, branch prediction units, decoders, execution units, reorder buffers, etc.). The caches **462A-462D** may comprise level 1 (L1) and level 2 (L2) caches. In addition, one or more shared caches **426** may be included in the caching hierarchy and shared by sets of the cores **460A-460D**. For example, one embodiment of the processor **407** includes 24 cores, each with its own L1 cache, twelve shared L2 caches, and twelve shared L3 caches. In this embodiment, one of the L2 and L3 caches are shared by two adjacent cores. The processor **407** and the graphics accelerator integration module **446** connect with system memory **441**, which may include processor memories **401-402**.

Coherency is maintained for data and instructions stored in the various caches **462A-462D**, **456** and system memory **441** via inter-core communication over a coherence bus **464**. For example, each cache may have cache coherency logic/circuitry associated therewith to communicate to over the coherence bus **464** in response to detected reads or writes to particular cache lines. In one implementation, a cache snooping protocol is implemented over the coherence bus **464** to snoop cache accesses. Cache snooping/coherency techniques are well understood by those of skill in the art and will not be described in detail here to avoid obscuring the underlying principles of the invention.

In one embodiment, a proxy circuit **425** communicatively couples the graphics acceleration module **446** to the coherence bus **464**, allowing the graphics acceleration module **446** to participate in the cache coherence protocol as a peer of the cores. In particular, an interface **435** provides connectivity to the proxy circuit **425** over high-speed link **440** (e.g., a PCIe bus, NVLink, etc.) and an interface **437** connects the graphics acceleration module **446** to the link **440**.

In one implementation, an accelerator integration circuit **436** provides cache management, memory access, context management, and interrupt management services on behalf of a plurality of graphics processing engines **431**, **432**, N of the graphics acceleration module **446**. The graphics processing engines **431**, **432**, N may each comprise a separate graphics processing unit (GPU). Alternatively, the graphics processing engines **431**, **432**, N may comprise different types of graphics processing engines within a GPU such as graphics execution units, media processing engines (e.g., video encoders/decoders), samplers, and blit engines. In other words, the graphics acceleration module may be a GPU with a plurality of graphics processing engines **431-432**, N or the graphics processing engines **431-432**, N may be individual GPUs integrated on a common package, line card, or chip.

In one embodiment, the accelerator integration circuit **436** includes a memory management unit (MMU) **439** for performing various memory management functions such as virtual-to-physical memory translations (also referred to as effective-to-real memory translations) and memory access protocols for accessing system memory **441**. The MMU **439** may also include a translation lookaside buffer (TLB) (not shown) for caching the virtual/effective to physical/real address translations. In one implementation, a cache **438** stores commands and data for efficient access by the graphics processing engines **431-432**, N. In one embodiment, the data stored in cache **438** and graphics memories **433-434**, N is kept coherent with the core caches **462A-462D**, **456** and system memory **411**. As mentioned, this may be accomplished via proxy circuit **425** which takes part in the cache coherency mechanism on behalf of cache **438** and memories **433-434**, N (e.g., sending updates to the cache **438** related to modifications/accesses of cache lines on processor caches **462A-462D**, **456** and receiving updates from the cache **438**).

A set of registers **445** store context data for threads executed by the graphics processing engines **431-432**, N and a context management circuit **448** manages the thread contexts. For example, the context management circuit **448** may perform save and restore operations to save and restore contexts of the various threads during contexts switches (e.g., where a first thread is saved and a second thread is stored so that the second thread can be executed by a graphics processing engine). For example, on a context switch, the context management circuit **448** may store current register values to a designated region in memory (e.g., identified by a context pointer). It may then restore the register values when returning to the context. In one embodiment, an interrupt management circuit **447** receives and processes interrupts received from system devices.

In one implementation, virtual/effective addresses from a graphics processing engine **431** are translated to real/physical addresses in system memory **411** by the MMU **439**. One embodiment of the accelerator integration circuit **436** supports multiple (e.g., 4, 8, 16) graphics accelerator modules **446** and/or other accelerator devices. The graphics accelerator module **446** may be dedicated to a single application executed on the processor **407** or may be shared between multiple applications. In one embodiment, a virtualized graphics execution environment is presented in which the resources of the graphics processing engines **431-432**, N are shared with multiple applications or virtual machines (VMs). The resources may be subdivided into “slices” which are allocated to different VMs and/or applications based on the processing requirements and priorities associated with the VMs and/or applications.

Thus, the accelerator integration circuit acts as a bridge to the system for the graphics acceleration module **446** and provides address translation and system memory cache services. In addition, the accelerator integration circuit **436** may provide virtualization facilities for the host processor to manage virtualization of the graphics processing engines, interrupts, and memory management.

Because hardware resources of the graphics processing engines **431-432**, N are mapped explicitly to the real address space seen by the host processor **407**, any host processor can address these resources directly using an effective address value. One function of the accelerator integration circuit **436**, in one embodiment, is the physical separation of the graphics processing engines **431-432**, N so that they appear to the system as independent units.

As mentioned, in the illustrated embodiment, one or more graphics memories **433-434**, M are coupled to each of the

graphics processing engines **431-432**, N, respectively. The graphics memories **433-434**, M store instructions and data being processed by each of the graphics processing engines **431-432**, N. The graphics memories **433-434**, M may be volatile memories such as DRAMs (including stacked DRAMs), GDDR memory (e.g., GDDR5, GDDR6), or HBM, and/or may be non-volatile memories such as 3D XPoint or Nano-Ram.

In one embodiment, to reduce data traffic over link **440**, biasing techniques are used to ensure that the data stored in graphics memories **433-434**, M is data which will be used most frequently by the graphics processing engines **431-432**, N and preferably not used by the cores **460A-460D** (at least not frequently). Similarly, the biasing mechanism attempts to keep data needed by the cores (and preferably not the graphics processing engines **431-432**, N) within the caches **462A-462D**, **456** of the cores and system memory **411**.

FIG. **4C** illustrates another embodiment in which the accelerator integration circuit **436** is integrated within the processor **407**. In this embodiment, the graphics processing engines **431-432**, N communicate directly over the high-speed link **440** to the accelerator integration circuit **436** via interface **437** and interface **435** (which, again, may be utilize any form of bus or interface protocol). The accelerator integration circuit **436** may perform the same operations as those described with respect to FIG. **4B**, but potentially at a higher throughput given its close proximity to the coherency bus **462** and caches **462A-462D**, **426**.

One embodiment supports different programming models including a dedicated-process programming model (no graphics acceleration module virtualization) and shared programming models (with virtualization). The latter may include programming models which are controlled by the accelerator integration circuit **436** and programming models which are controlled by the graphics acceleration module **446**.

In one embodiment of the dedicated process model, graphics processing engines **431-432**, N are dedicated to a single application or process under a single operating system. The single application can funnel other application requests to the graphics engines **431-432**, N, providing virtualization within a VM/partition.

In the dedicated-process programming models, the graphics processing engines **431-432**, N, may be shared by multiple VM/application partitions. The shared models require a system hypervisor to virtualize the graphics processing engines **431-432**, N to allow access by each operating system. For single-partition systems without a hypervisor, the graphics processing engines **431-432**, N are owned by the operating system. In both cases, the operating system can virtualize the graphics processing engines **431-432**, N to provide access to each process or application.

For the shared programming model, the graphics acceleration module **446** or an individual graphics processing engine **431-432**, N selects a process element using a process handle. In one embodiment, process elements are stored in system memory **411** and are addressable using the effective address to real address translation techniques described herein. The process handle may be an implementation-specific value provided to the host process when registering its context with the graphics processing engine **431-432**, N (that is, calling system software to add the process element to the process element linked list). The lower 16-bits of the process handle may be the offset of the process element within the process element linked list.

FIG. **4D** illustrates an exemplary accelerator integration slice **490**. As used herein, a “slice” comprises a specified

portion of the processing resources of the accelerator integration circuit **436**. Application effective address space **482** within system memory **411** stores process elements **483**. In one embodiment, the process elements **483** are stored in response to GPU invocations **481** from applications **480** executed on the processor **407**. A process element **483** contains the process state for the corresponding application **480**. A work descriptor (WD) **484** contained in the process element **483** can be a single job requested by an application or may contain a pointer to a queue of jobs. In the latter case, the WD **484** is a pointer to the job request queue in the application's address space **482**.

The graphics acceleration module **446** and/or the individual graphics processing engines **431-432, N** can be shared by all or a subset of the processes in the system. Embodiments of the invention include an infrastructure for setting up the process state and sending a WD **484** to a graphics acceleration module **446** to start a job in a virtualized environment.

In one implementation, the dedicated-process programming model is implementation-specific. In this model, a single process owns the graphics acceleration module **446** or an individual graphics processing engine **431**. Because the graphics acceleration module **446** is owned by a single process, the hypervisor initializes the accelerator integration circuit **436** for the owning partition and the operating system initializes the accelerator integration circuit **436** for the owning process at the time when the graphics acceleration module **446** is assigned.

In operation, a WD fetch unit **491** in the accelerator integration slice **490** fetches the next WD **484** which includes an indication of the work to be done by one of the graphics processing engines of the graphics acceleration module **446**. Data from the WD **484** may be stored in registers **445** and used by the MMU **439**, interrupt management circuit **447** and/or context management circuit **446** as illustrated. For example, one embodiment of the MMU **439** includes segment/page walk circuitry for accessing segment/page tables **486** within the OS virtual address space **485**. The interrupt management circuit **447** may process interrupt events **492** received from the graphics acceleration module **446**. When performing graphics operations, an effective address **493** generated by a graphics processing engine **431-432, N** is translated to a real address by the MMU **439**.

In one embodiment, the same set of registers **445** are duplicated for each graphics processing engine **431-432, N** and/or graphics acceleration module **446** and may be initialized by the hypervisor or operating system. Each of these duplicated registers may be included in an accelerator integration slice **490**. Exemplary registers that may be initialized by the hypervisor are shown in Table 1.

TABLE 1

Hypervisor Initialized Registers	
1	Slice Control Register
2	Real Address (RA) Scheduled Processes Area Pointer
3	Authority Mask Override Register
4	Interrupt Vector Table Entry Offset
5	Interrupt Vector Table Entry Limit
6	State Register
7	Logical Partition ID
8	Real address (RA) Hypervisor Accelerator Utilization Record Pointer
9	Storage Description Register

Exemplary registers that may be initialized by the operating system are shown in Table 2.

TABLE 2

Operating System Initialized Registers	
1	Process and Thread Identification
2	Effective Address (EA) Context Save/Restore Pointer
3	Virtual Address (VA) Accelerator Utilization Record Pointer
4	Virtual Address (VA) Storage Segment Table Pointer
5	Authority Mask
6	Work descriptor

In one embodiment, each WD **484** is specific to a particular graphics acceleration module **446** and/or graphics processing engines **431-432, N**. It contains all the information a graphics processing engine **431-432, N** requires to do its work or it can be a pointer to a memory location where the application has set up a command queue of work to be completed.

FIG. 4E illustrates additional details for one embodiment of a shared model. This embodiment includes a hypervisor real address space **498** in which a process element list **499** is stored. The hypervisor real address space **498** is accessible via a hypervisor **496** which virtualizes the graphics acceleration module engines for the operating system **495**.

The shared programming models allow for all or a subset of processes from all or a subset of partitions in the system to use a graphics acceleration module **446**. There are two programming models where the graphics acceleration module **446** is shared by multiple processes and partitions: time-sliced shared and graphics directed shared.

In this model, the system hypervisor **496** owns the graphics acceleration module **446** and makes its function available to all operating systems **495**. For a graphics acceleration module **446** to support virtualization by the system hypervisor **496**, the graphics acceleration module **446** may adhere to the following requirements: 1) An application's job request must be autonomous (that is, the state does not need to be maintained between jobs), or the graphics acceleration module **446** must provide a context save and restore mechanism. 2) An application's job request is guaranteed by the graphics acceleration module **446** to complete in a specified amount of time, including any translation faults, or the graphics acceleration module **446** provides the ability to preempt the processing of the job. 3) The graphics acceleration module **446** must be guaranteed fairness between processes when operating in the directed shared programming model.

In one embodiment, for the shared model, the application **480** is required to make an operating system **495** system call with a graphics acceleration module **446** type, a work descriptor (WD), an authority mask register (AMR) value, and a context save/restore area pointer (CSR). The graphics acceleration module **446** type describes the targeted acceleration function for the system call. The graphics acceleration module **446** type may be a system-specific value. The WD is formatted specifically for the graphics acceleration module **446** and can be in the form of a graphics acceleration module **446** command, an effective address pointer to a user-defined structure, an effective address pointer to a queue of commands, or any other data structure to describe the work to be done by the graphics acceleration module **446**. In one embodiment, the AMR value is the AMR state to use for the current process. The value passed to the operating system is similar to an application setting the AMR. If the accelerator integration circuit **436** and graphics acceleration module **446** implementations do not support a User Authority Mask Override Register (UAMOR), the operating system may apply the current UAMOR value to

the AMR value before passing the AMR in the hypervisor call. The hypervisor **496** may optionally apply the current Authority Mask Override Register (AMOR) value before placing the AMR into the process element **483**. In one embodiment, the CSRP is one of the registers **445** containing the effective address of an area in the application's address space **482** for the graphics acceleration module **446** to save and restore the context state. This pointer is optional if no state is required to be saved between jobs or when a job is preempted. The context save/restore area may be pinned system memory.

Upon receiving the system call, the operating system **495** may verify that the application **480** has registered and been given the authority to use the graphics acceleration module **446**. The operating system **495** then calls the hypervisor **496** with the information shown in Table 3.

TABLE 3

OS to Hypervisor Call Parameters	
1	A work descriptor (WD)
2	An Authority Mask Register (AMR) value (potentially masked).
3	An effective address (EA) Context Save/Restore Area Pointer (CSRP)
4	A process ID (PID) and optional thread ID (TID)
5	A virtual address (VA) accelerator utilization record pointer (AURP)
6	The virtual address of the storage segment table pointer (SSTP)
7	A logical interrupt service number (LISN)

Upon receiving the hypervisor call, the hypervisor **496** verifies that the operating system **495** has registered and been given the authority to use the graphics acceleration module **446**. The hypervisor **496** then puts the process element **483** into the process element linked list for the corresponding graphics acceleration module **446** type. The process element may include the information shown in Table 4

TABLE 4

Process Element Information	
1	A work descriptor (WD)
2	An Authority Mask Register (AMR) value (potentially masked).
3	An effective address (EA) Context Save/Restore Area Pointer (CSRP)
4	A process ID (PID) and optional thread ID (TID)
5	A virtual address (VA) accelerator utilization record pointer (AURP)
6	The virtual address of the storage segment table pointer (SSTP)
7	A logical interrupt service number (LISN)
8	Interrupt vector table, derived from the hypervisor call parameters.
9	A state register (SR) value
10	A logical partition ID (LPID)
11	A real address (RA) hypervisor accelerator utilization record pointer
12	The Storage Descriptor Register (SDR)

In one embodiment, the hypervisor initializes a plurality of accelerator integration slice **490** registers **445**.

As illustrated in FIG. 4F, one embodiment of the invention employs a unified memory addressable via a common virtual memory address space used to access the physical processor memories **401-402** and GPU memories **420-423**. In this implementation, operations executed on the GPUs **410-413** utilize the same virtual/effective memory address space to access the processors memories **401-402** and vice versa, thereby simplifying programmability. In one embodiment, a first portion of the virtual/effective address space is allocated to the processor memory **401**, a second portion to the second processor memory **402**, a third portion to the GPU memory **420**, and so on. The entire virtual/effective

memory space (sometimes referred to as the effective address space) is thereby distributed across each of the processor memories **401-402** and GPU memories **420-423**, allowing any processor or GPU to access any physical memory with a virtual address mapped to that memory.

In one embodiment, bias/coherence management circuitry **494A-494E** within one or more of the MMUs **439A-439E** ensures cache coherence between the caches of the host processors (e.g., **405**) and the GPUs **410-413** and implements biasing techniques indicating the physical memories in which certain types of data should be stored. While multiple instances of bias/coherence management circuitry **494A-494E** are illustrated in FIG. 4F, the bias/coherence circuitry may be implemented within the MMU of one or more host processors **405** and/or within the accelerator integration circuit **436**.

One embodiment allows GPU-attached memory **420-423** to be mapped as part of system memory, and accessed using shared virtual memory (SVM) technology, but without suffering the typical performance drawbacks associated with full system cache coherence. The ability to GPU-attached memory **420-423** to be accessed as system memory without onerous cache coherence overhead provides a beneficial operating environment for GPU offload. This arrangement allows the host processor **405** software to setup operands and access computation results, without the overhead of traditional I/O DMA data copies. Such traditional copies involve driver calls, interrupts and memory mapped I/O (MMIO) accesses that are all inefficient relative to simple memory accesses. At the same time, the ability to access GPU attached memory **420-423** without cache coherence overheads can be critical to the execution time of an off-loaded computation. In cases with substantial streaming write memory traffic, for example, cache coherence overhead can significantly reduce the effective write bandwidth seen by a GPU **410-413**. The efficiency of operand setup, the efficiency of results access, and the efficiency of GPU computation all play a role in determining the effectiveness of GPU offload.

In one implementation, the selection of between GPU bias and host processor bias is driven by a bias tracker data structure. A bias table may be used, for example, which may be a page-granular structure (i.e., controlled at the granularity of a memory page) that includes 1 or 2 bits per GPU-attached memory page. The bias table may be implemented in a stolen memory range of one or more GPU-attached memories **420-423**, with or without a bias cache in the GPU **410-413** (e.g., to cache frequently/recently used entries of the bias table). Alternatively, the entire bias table may be maintained within the GPU.

In one implementation, the bias table entry associated with each access to the GPU-attached memory **420-423** is accessed prior the actual access to the GPU memory, causing the following operations. First, local requests from the GPU **410-413** that find their page in GPU bias are forwarded directly to a corresponding GPU memory **420-423**. Local requests from the GPU that find their page in host bias are forwarded to the processor **405** (e.g., over a high-speed link as discussed above). In one embodiment, requests from the processor **405** that find the requested page in host processor bias complete the request like a normal memory read. Alternatively, requests directed to a GPU-biased page may be forwarded to the GPU **410-413**. The GPU may then transition the page to a host processor bias if it is not currently using the page.

The bias state of a page can be changed either by a software-based mechanism, a hardware-assisted software-based mechanism, or, for a limited set of cases, a purely hardware-based mechanism.

One mechanism for changing the bias state employs an API call (e.g. OpenCL), which, in turn, calls the GPU's device driver which, in turn, sends a message (or enqueues a command descriptor) to the GPU directing it to change the bias state and, for some transitions, perform a cache flushing operation in the host. The cache flushing operation is required for a transition from host processor **405** bias to GPU bias, but is not required for the opposite transition.

In one embodiment, cache coherency is maintained by temporarily rendering GPU-biased pages uncacheable by the host processor **405**. To access these pages, the processor **405** may request access from the GPU **410** which may or may not grant access right away, depending on the implementation. Thus, to reduce communication between the processor **405** and GPU **410** it is beneficial to ensure that GPU-biased pages are those which are required by the GPU but not the host processor **405** and vice versa.

Graphics Processing Pipeline

FIG. 5 illustrates a graphics processing pipeline **500**, according to an embodiment. In one embodiment, a graphics processor can implement the illustrated graphics processing pipeline **500**. The graphics processor can be included within the parallel processing subsystems as described herein, such as the parallel processor **200** of FIG. 2A, which, in one embodiment, is a variant of the parallel processor(s) **112** of FIG. 1. The various parallel processing systems can implement the graphics processing pipeline **500** via one or more instances of the parallel processing unit (e.g., parallel processing unit **202** of FIG. 2A) as described herein. For example, a shader unit (e.g., graphics multiprocessor **234** of FIG. 2D) may be configured to perform the functions of one or more of a vertex processing unit **504**, a tessellation control processing unit **508**, a tessellation evaluation processing unit **512**, a geometry processing unit **516**, and a fragment/pixel processing unit **524**. The functions of data assembler **502**, primitive assemblers **506**, **514**, **518**, tessellation unit **510**, rasterizer **522**, and raster operations unit **526** may also be performed by other processing engines within a processing cluster (e.g., processing cluster **214** of FIG. 3A) and a corresponding partition unit (e.g., partition unit **220A-220N** of FIG. 2C). The graphics processing pipeline **500** may also be implemented using dedicated processing units for one or more functions. In one embodiment, one or more portions of the graphics processing pipeline **500** can be performed by parallel processing logic within a general-purpose processor (e.g., CPU). In one embodiment, one or more portions of the graphics processing pipeline **500** can access on-chip memory (e.g., parallel processor memory **222** as in FIG. 2A) via a memory interface **528**, which may be an instance of the memory interface **218** of FIG. 2A.

In one embodiment, the data assembler **502** is a processing unit that collects vertex data for surfaces and primitives. The data assembler **502** then outputs the vertex data, including the vertex attributes, to the vertex processing unit **504**. The vertex processing unit **504** is a programmable execution unit that executes vertex shader programs, lighting and transforming vertex data as specified by the vertex shader programs. The vertex processing unit **504** reads data that is stored in cache, local or system memory for use in processing the vertex data and may be programmed to transform the vertex data from an object-based coordinate representation to a world space coordinate space or a normalized device coordinates space.

A first instance of a primitive assembler **506** receives vertex attributes from the vertex processing unit **504**. The primitive assembler **506** reads stored vertex attributes as needed and constructs graphics primitives for processing by tessellation control processing unit **508**. The graphics primitives include triangles, line segments, points, patches, and so forth, as supported by various graphics processing application programming interfaces (APIs).

The tessellation control processing unit **508** treats the input vertices as control points for a geometric patch. The control points are transformed from an input representation from the patch (e.g., the patch's bases) to a representation that is suitable for use in surface evaluation by the tessellation evaluation processing unit **512**. The tessellation control processing unit **508** can also compute tessellation factors for edges of geometric patches. A tessellation factor applies to a single edge and quantifies a view-dependent level of detail associated with the edge. A tessellation unit **510** is configured to receive the tessellation factors for edges of a patch and to tessellate the patch into multiple geometric primitives such as line, triangle, or quadrilateral primitives, which are transmitted to a tessellation evaluation processing unit **512**. The tessellation evaluation processing unit **512** operates on parameterized coordinates of the subdivided patch to generate a surface representation and vertex attributes for each vertex associated with the geometric primitives.

A second instance of a primitive assembler **514** receives vertex attributes from the tessellation evaluation processing unit **512**, reading stored vertex attributes as needed, and constructs graphics primitives for processing by the geometry processing unit **516**. The geometry processing unit **516** is a programmable execution unit that executes geometry shader programs to transform graphics primitives received from primitive assembler **514** as specified by the geometry shader programs. In one embodiment, the geometry processing unit **516** is programmed to subdivide the graphics primitives into one or more new graphics primitives and calculate parameters used to rasterize the new graphics primitives.

In some embodiments, the geometry processing unit **516** can add or delete elements in the geometry stream. The geometry processing unit **516** outputs the parameters and vertices specifying new graphics primitives to primitive assembler **518**. The primitive assembler **518** receives the parameters and vertices from the geometry processing unit **516** and constructs graphics primitives for processing by a viewport scale, cull, and clip unit **520**. The geometry processing unit **516** reads data that is stored in parallel processor memory or system memory for use in processing the geometry data. The viewport scale, cull, and clip unit **520** performs clipping, culling, and viewport scaling and outputs processed graphics primitives to a rasterizer **522**.

The rasterizer **522** can perform depth culling and other depth-based optimizations. The rasterizer **522** also performs scan conversion on the new graphics primitives to generate fragments and outputs those fragments and associated coverage data to the fragment/pixel processing unit **524**.

The fragment/pixel processing unit **524** is a programmable execution unit that is configured to execute fragment shader programs or pixel shader programs. The fragment/pixel processing unit **524** transforming fragments or pixels received from rasterizer **522**, as specified by the fragment or pixel shader programs. For example, the fragment/pixel processing unit **524** may be programmed to perform operations included but not limited to texture mapping, shading, blending, texture correction and perspective correction to

produce shaded fragments or pixels that are output to a raster operations unit **526**. The fragment/pixel processing unit **524** can read data that is stored in either the parallel processor memory or the system memory for use when processing the fragment data. Fragment or pixel shader programs may be configured to shade at sample, pixel, tile, or other granularities, depending on the sampling rate configured for the processing units.

The raster operations unit **526** is a processing unit that performs raster operations including, but not limited to stencil, z test, blending, and the like, and outputs pixel data as processed graphics data to be stored in graphics memory, e.g., parallel processor memory **222** as in FIG. 2A, and/or system memory **104** as in FIG. 1, to be displayed on the one or more display device(s) **110** or for further processing by one of the one or more processor(s) **102** or parallel processor(s) **112**. In some embodiments, the raster operations unit **526** is configured to compress z or color data that is written to memory and decompress z or color data that is read from memory.

FIG. 6 illustrates a computing device **600** hosting a dynamic atomic merging mechanism (“merging mechanism”) **610** according to one embodiment. Computing device **600** represents a communication and data processing device including (but not limited to) smart wearable devices, smartphones, virtual reality (VR) devices, head-mounted display (HMDs), mobile computers, Internet of Things (IoT) devices, laptop computers, desktop computers, server computers, etc., and be similar to or the same as computing device **100** of FIG. 1; accordingly, for brevity, clarity, and ease of understanding, many of the details stated above with reference to FIGS. 1-5 are not further discussed or repeated hereafter.

Computing device **600** may further include (without limitations) an autonomous machine or an artificially intelligent agent, such as a mechanical agent or machine, an electronics agent or machine, a virtual agent or machine, an electro-mechanical agent or machine, etc. Examples of autonomous machines or artificially intelligent agents may include (without limitation) robots, autonomous vehicles (e.g., self-driving cars, self-flying planes, self-sailing boats, etc.), autonomous equipment (self-operating construction vehicles, self-operating medical equipment, etc.), and/or the like. Throughout this document, “computing device” may be interchangeably referred to as “autonomous machine” or “artificially intelligent agent” or simply “robot”.

Computing device **600** may further include (without limitations) large computing systems, such as server computers, desktop computers, etc., and may further include set-top boxes (e.g., Internet-based cable television set-top boxes, etc.), global positioning system (GPS)-based devices, etc. Computing device **600** may include mobile computing devices serving as communication devices, such as cellular phones including smartphones, personal digital assistants (PDAs), tablet computers, laptop computers, e-readers, smart televisions, television platforms, wearable devices (e.g., glasses, watches, bracelets, smartcards, jewelry, clothing items, etc.), media players, etc. For example, in one embodiment, computing device **600** may include a mobile computing device employing a computer platform hosting an integrated circuit (“IC”), such as system on a chip (“SoC” or “SOC”), integrating various hardware and/or software components of computing device **600** on a single chip.

As illustrated, in one embodiment, computing device **600** may include any number and type of hardware and/or software components, such as (without limitation) graphics processing unit (“GPU” or simply “graphics processor”)

614, graphics driver (also referred to as “GPU driver”, “graphics driver logic”, “driver logic”, user-mode driver (UMD), UMD, user-mode driver framework (UMDF), UMDF, or simply “driver”) **616**, central processing unit (“CPU” or simply “application processor”) **612**, memory **608**, network devices, drivers, or the like, as well as input/output (I/O) sources **604**, such as touchscreens, touch panels, touch pads, virtual or regular keyboards, virtual or regular mice, ports, connectors, etc. Computing device **600** may include operating system (OS) **606** serving as an interface between hardware and/or physical resources of the computer device **600** and a user. It is contemplated that graphics processor **614** and application processor **612** may be one or more of processor(s) **102** of FIG. 1.

It is to be appreciated that a lesser or more equipped system than the example described above may be preferred for certain implementations. Therefore, the configuration of computing device **600** may vary from implementation to implementation depending upon numerous factors, such as price constraints, performance requirements, technological improvements, or other circumstances.

Embodiments may be implemented as any or a combination of: one or more microchips or integrated circuits interconnected using a parentboard, hardwired logic, software stored by a memory device and executed by a microprocessor, firmware, an application specific integrated circuit (ASIC), and/or a field programmable gate array (FPGA). The terms “logic”, “module”, “component”, “engine”, and “mechanism” may include, by way of example, software or hardware and/or combinations of software and hardware.

In one embodiment, merging mechanism **610** may be hosted or facilitated by operating system **606** of computing device **600**. In another embodiment, merging mechanism **610** may be hosted by or part of graphics processing unit (“GPU” or simply “graphics processor”) **614** or firmware of graphics processor **614**. For example, merging mechanism **610** may be embedded in or implemented as part of the processing hardware of graphics processor **614**. Similarly, in yet another embodiment, merging mechanism **610** may be hosted by or part of central processing unit (“CPU” or simply “application processor”) **612**. For example, merging mechanism **610** may be embedded in or implemented as part of the processing hardware of application processor **612**. In yet another embodiment, merging mechanism **610** may be hosted by or part of any number and type of components of computing device **600**, such as a portion of merging mechanism **610** may be hosted by or part of operating system **606**, another portion may be hosted by or part of graphics processor **614**, another portion may be hosted by or part of application processor **612**, while one or more portions of merging mechanism **610** may be hosted by or part of operating system **606** and/or any number and type of devices of computing device **600**. It is contemplated that one or more portions or components of merging mechanism **610** may be employed as hardware, software, and/or firmware.

It is contemplated that embodiments are not limited to any particular implementation or hosting of merging mechanism **610** and that merging mechanism **610** and one or more of its components may be implemented as hardware, software, firmware, or any combination thereof.

Computing device **600** may host network interface(s) to provide access to a network, such as a LAN, a wide area network (WAN), a metropolitan area network (MAN), a personal area network (PAN), Bluetooth, a cloud network, a mobile network (e.g., 3rd Generation (3G), 4th Generation

(4G), etc.), an intranet, the Internet, etc. Network interface(s) may include, for example, a wireless network interface having antenna, which may represent one or more antenna(e). Network interface(s) may also include, for example, a wired network interface to communicate with remote devices via network cable, which may be, for example, an Ethernet cable, a coaxial cable, a fiber optic cable, a serial cable, or a parallel cable.

Embodiments may be provided, for example, as a computer program product which may include one or more machine-readable media having stored thereon machine-executable instructions that, when executed by one or more machines such as a computer, network of computers, or other electronic devices, may result in the one or more machines carrying out operations in accordance with embodiments described herein. A machine-readable medium may include, but is not limited to, floppy diskettes, optical disks, CD-ROMs (Compact Disc-Read Only Memories), and magneto-optical disks, ROMs, RAMs, EPROMs (Erasable Programmable Read Only Memories), EEPROMs (Electrically Erasable Programmable Read Only Memories), magnetic or optical cards, flash memory, or other type of media/machine-readable medium suitable for storing machine-executable instructions.

Moreover, embodiments may be downloaded as a computer program product, wherein the program may be transferred from a remote computer (e.g., a server) to a requesting computer (e.g., a client) by way of one or more data signals embodied in and/or modulated by a carrier wave or other propagation medium via a communication link (e.g., a modem and/or network connection).

Throughout the document, term “user” may be interchangeably referred to as “viewer”, “observer”, “person”, “individual”, “end-user”, and/or the like. It is to be noted that throughout this document, terms like “graphics domain” may be referenced interchangeably with “graphics processing unit”, “graphics processor”, or simply “GPU” and similarly, “CPU domain” or “host domain” may be referenced interchangeably with “computer processing unit”, “application processor”, or simply “CPU”.

It is to be noted that terms like “node”, “computing node”, “server”, “server device”, “cloud computer”, “cloud server”, “cloud server computer”, “machine”, “host machine”, “device”, “computing device”, “computer”, “computing system”, and the like, may be used interchangeably throughout this document. It is to be further noted that terms like “application”, “software application”, “program”, “software program”, “package”, “software package”, and the like, may be used interchangeably throughout this document. Also, terms like “job”, “input”, “request”, “message”, and the like, may be used interchangeably throughout this document.

FIG. 7 illustrates merging mechanism 610 of FIG. 6 according to one embodiment. For brevity, many of the details already discussed with reference to FIGS. 1-6 are not repeated or discussed hereafter. In one embodiment, merging mechanism 610 may include any number and type of components, such as (without limitations): detection/comparison logic 701; merging/conversion logic 703; generation/execution logic 705; and communication/compatibility logic 707.

From three-dimensional (3D) and compute workload analysis, there are several cases where multiple slots of an atomic SIMD message have the same address. For example, this is fairly common in “histogram” code, where a histogram bucket is incremented based on a pixel’s grayscale value. Since the slots in an atomic SIMD message process adjacent pixels, there is a high probability that the pixels are

likely to have the same or close enough grayscale values and thus the slots will increment the same histogram bucket, such as the same Dword address (32 bit atomic) or Word address (16 bit atomic).

In existing GDP or SLM controller design, when multiple SIMD slots of an atomic message map to the same address, the atomic operations are serialized (such as processed one per cycle), which, in turn, leads to an atomic SIMD16 message (16 slots) taking up to 16 cycles to complete all its atomic operations

Embodiments provide for a novel technique for merging multiple atomic operations to memory into a single atomic operation. In one embodiment, this novel and innovative technique may be implemented in hardware and used to detect such cases and merge any multiple same-address atomic operations in to a single atomic operation. Thus, an atomic SIMD16 operation may be completed in a single cycle, resulting in significantly speeding up these applications. This idea can be applied to both SLM and global memory systems.

Typically, for example, a graphics shader execution unit (EU) may send a number of atomic SIMD messages to a data port controller, where this data port controller computes a memory address per slot and sends the atomic operation to a cache controller with the atomic arithmetic logic unit (ALU). However, in existing techniques, when the data port detects that two or more of the slot addresses are the same, it simply serialized those slots or, in other words, sends one slot atomic operation per cycle to the cache controller. This is primarily due to data dependency between these atomic operations (e.g., a current or younger atomic operation will need the result of a previous or older atomic operation to the same address) that can cause long delays in processing of atomic operations.

In one embodiment, detection/comparison logic 701 of merging mechanism 610 may be used to monitor atomic messages and detect any slot addresses and then compare the slot addresses to detect any address match. Upon detection of an address match, merging/conversion logic 703 may then be triggered to map the matching slot addresses and create a two-dimensional (2D) merge signal vector to indicate which of the slots are needed to merged into one or more groups. In one embodiment, generation/execution logic 705 may then be used to generate one or more new merged atomic operations for all the merge groups.

For example, as will be further described and illustrate with reference to FIG. 9, an SIMG atomic message may be issued by an EU, passing through a number of slots to then be processed by merging mechanism 610. In one embodiment, using one or more components of merging mechanism 610, unique addresses are retained, while other atomic operations are merged, resulting in a far fewer number of atomic operations.

Further, communication/compatibility logic 707 may be used to facilitate the needed communication and compatibility between any number of devices of computing device 600 and as well as with and between various components of merging mechanism 610.

Communication/compatibility logic 707 may be used to facilitate dynamic communication and compatibility between computing device 600 and any number and type of other computing devices (such as mobile computing device, desktop computer, server computing device, etc.); processing devices or components (such as CPUs, GPUs, etc.); capturing/sensing/detecting devices (such as capturing/sensing components including cameras, depth sensing cameras, camera sensors, red green blue (“RGB” or “rgb”) sensors,

microphones, etc.); display devices (such as output components including display screens, display areas, display projectors, etc.); user/context-awareness components and/or identification/verification sensors/devices (such as biometric sensors/detectors, scanners, etc.); database(s) **730**, such as memory or storage devices, databases, and/or data sources (such as data storage devices, hard drives, solid-state drives, hard disks, memory cards or devices, memory circuits, etc.); communication medium(s) **725**, such as one or more communication channels or networks (e.g., cloud networks, the Internet, intranets, cellular networks, proximity networks, such as Bluetooth, Bluetooth low energy (BLE), Bluetooth Smart, Wi-Fi proximity, Radio Frequency Identification (RFID), Near Field Communication (NFC), Body Area Network (BAN), etc.); wireless or wired communications and relevant protocols (e.g., Wi-Fi®, WiMAX, Ethernet, etc.); connectivity and location management techniques; software applications/websites (e.g., social and/or business networking websites, etc., business applications, games and other entertainment applications, etc.); and programming languages, etc., while ensuring compatibility with changing technologies, parameters, protocols, standards, etc.

Throughout this document, terms like “logic”, “component”, “module”, “framework”, “engine”, “mechanism”, and the like, may be referenced interchangeably and include, by way of example, software, hardware, firmware, and/or any combination thereof. In one example, “logic” may refer to or include a software component that is capable of working with one or more of an operating system (e.g., operating system **606**), a graphics driver (e.g., graphics driver **616**), etc., of a computing device, such as computing device **600**. In another example, “logic” may refer to or include a hardware component that is capable of being physically installed along with or as part of one or more system hardware elements, such as an application processor (e.g., CPU **612**), a graphics processor (e.g., GPU **614**), etc., of a computing device, such as computing device **600**. In yet another embodiment, “logic” may refer to or include a firmware component that is capable of being part of system firmware, such as firmware of an application processor (e.g., CPU **612**) or a graphics processor (e.g., GPU **614**), etc., of a computing device, such as computing device **600**.

Further, any use of a particular brand, word, term, phrase, name, and/or acronym, such as “monitoring”, “detecting”, “mapping”, “creating”, “generating”, “merging”, “atomic operation”, “SIMD”, “EU”, “client”, “graphics cache”, “GPU”, “graphics processor”, “GPU domain”, “GPGPU”, “CPU”, “application processor”, “CPU domain”, “graphics driver”, “workload”, “application”, “graphics pipeline”, “pipeline processes”, “API”, “3D API”, “OpenGL®”, “DirectX®”, “hardware”, “software”, “agent”, “graphics driver”, “kernel mode graphics driver”, “user-mode driver”, “user-mode driver framework”, “buffer”, “graphics buffer”, “task”, “process”, “operation”, “software application”, “game”, etc., should not be read to limit embodiments to software or devices that carry that label in products or in literature external to this document.

It is contemplated that any number and type of components may be added to and/or removed from merging mechanism **610** to facilitate various embodiments including adding, removing, and/or enhancing certain features. For brevity, clarity, and ease of understanding of merging mechanism **610**, many of the standard and/or known components, such as those of a computing device, are not shown or discussed here. It is contemplated that embodiments, as described herein, are not limited to any particular technol-

ogy, topology, system, architecture, and/or standard and are dynamic enough to adopt and adapt to any future changes.

FIG. **8** illustrates a conventional atomic message pipeline **800**. As illustrated, pipeline **800** begins with graphics shader **801** issuing or sending atomic SIMD messages corresponding to slot memory addresses **803** to data port controller **805** for processing, where data port controller **805** computes memory addresses for each slot address **803** and sends atomic operations **807** to cache controller (with atomic ALU) **809** for further processing.

As illustrated, in this conventional pipeline **800**, when data port controller **805** detects that two or more slot addresses **803** are the same, it serializes these slots, such as sending each slot atomic operation **807** per cycle to cache controller **809**. This is primarily because there are data dependencies between these atomic operations **807** (such a younger atomic operation needing or relying upon results from an older atomic operation to the same address), which can cause long delays in processing atomic operations **807**.

FIG. **9A** illustrates a novel atomic message pipeline **900** allowing for dynamic merging of atomic operations according to one embodiment. For brevity, many of the details previously discussed with reference to FIGS. **1-8** may not be discussed or repeated hereafter. Any processes relating to pipeline **900** may be performed by processing logic that may comprise hardware (e.g., circuitry, dedicated logic, programmable logic, etc.), software (such as instructions run on a processing device), or a combination thereof, as facilitated by merging mechanism **900** of FIG. **6**. The processes associated with pipeline **900** may be illustrated or recited in linear sequences for brevity and clarity in presentation; however, it is contemplated that any number of them can be performed in parallel, asynchronously, or in different orders.

In the illustrated embodiment, as described with reference to FIG. **7**, merging mechanism **610** at controller **905** (e.g., data port controller, SLM controller, etc.) is used for comparison of slot addresses **907** to detect any matches between slot addresses **907** as facilitated by detection/comparison logic **701** of FIG. **7** and subsequently, create a 2D merging of signal vector which indicates which slot addresses **907** are to be merged into one or more groups as facilitated by merging/conversion logic **703** of FIG. **7**. Further, in one embodiment, merging mechanism **610** to generate one or more new merged atomic operations **907** for all the merged groups as facilitated by generation/execution logic **705** of FIG. **7**.

As illustrated, EU **901** (e.g., graphics shader EU) issues an atomic message, such as SIMD8 atomic message, where the first 2 slot addresses **911**, such as Slot0 Atomic_inc(A) and Slot1 Atomic_inc(A), of slot addresses **903** operate on address A, next 2 slot addresses **913**, such as Slot2 Atomic_inc(B) and Slot3 Atomic_inc(B), of slot addresses **903** operate on address B, while the last four slot addresses **915**, such as Slot4 Atomic_inc(C), Slot5 Atomic_inc(C), Slot6 Atomic_inc(C), and Slot7 Atomic_inc(C), of slot addresses **903** operate on address C.

In one embodiment, detection/comparison **701** of FIG. **7** detects all the unique addresses of slot addresses **903** and the number of instances of each of those unique addresses. Then, merging/conversion logic **703** of FIG. **7** may be triggered to convert the atomic_inc to atomic_add operation, 1 per unique address. In this illustrated embodiment, the output of merging mechanism **610** offers three atomic-add operations **907**, such as one each for addresses A **921**, B **923**, and C **925**. A **921** and B **923** are incremented by 2, while C **925** is incremented by 4, reducing the total number to merely 3 as opposed to 8.

The algorithm recited below may be used for Atomic_Increment and Atomic_Decrement messages without any data return and this technique may also be extended to other atomic math operations. For example, merging mechanism **610** may be used to count how many slot addresses **903** of an ATOMIC INC/DEC/PREDEC message are matching on the same address. The signal array slot merge[i][j] is generated to indicate if the dword address of slot[i] is same as the dword address of slot[j], and in that case slot[j] is merged with slot[i] (i<j).

For example, FIG. 9B illustrates this array **950** for 32 slots, where this is an upper triangular matrix and this signal array may be used to detect and fuse atomics conflicting on the same address. For brevity, many of the details previously discussed with reference to FIGS. 1-9A may not be discussed or repeated hereafter.

Next, to “fuse” atomics that are conflicting on the same address, in one embodiment, PopCount logic may implemented as part of merging/conversion logic **703** of merging mechanism **610** for every row of this array, where PopCount logic counts the number of set bits in the row. In one embodiment, the following algorithm may be implemented in one or more of SLM and data port controllers for handling local memory and global memory, respectively. For example, in one implementation, certain multiplexers may be needed to replace the atomic opcode and atomic operand per slot. The algorithm to fuse atomics is shown below: for every row i:

```

n := popCount(slot_merge[i][ ]); //count the # of conflicts
on dword or word address(i)
if (n > 0) then
  if (message_opcode == ATOMIC_INC) then
    replace the opcode of slot[i]with ATOMIC_ADD(n+1)
    drop all slots j for which slot_merge[i][j] is 1
  else if (message_opcode == ATOMIC_DEC or
  ATOMIC_PREDEC) then
    replace the opcode of slot[i] with ATOMIC_SUB(n+1)
    drop all slots j for which slot_merge[i][j] is 1
  end if
end if
end for

```

FIG. 9C illustrates a method **970** for merging of atomic operations for memory according to one embodiment. For brevity, many of the details previously discussed with reference to FIGS. 1-9B may not be discussed or repeated hereafter. Any processes relating to method **970** may be performed by processing logic that may comprise hardware (e.g., circuitry, dedicated logic, programmable logic, etc.), software (such as instructions run on a processing device), or a combination thereof, as facilitated by merging mechanism **610** of FIG. 6. The processes associated with method **970** may be illustrated or recited in linear sequences for brevity and clarity in presentation; however, it is contemplated that any number of them can be performed in parallel, asynchronously, or in different orders.

Method **970** begins at block **971** with detection of an atomic message and slot addresses. At block **973**, the slot addresses are compared with each other to detect any addresses that match with any of the other addresses. At block **975**, the matched addresses are merged into one or more merged groups. At block **977**, one or more new merged atomic operations are generated corresponding to the one or more merged groups.

This novel technique has been described above with reference to specific embodiments. Persons skilled in the art, however, will understand that various modifications and

changes may be made thereto without departing from the broader spirit and scope as set forth in the appended claims. The foregoing description and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

5 System Overview II

FIG. 10 is a block diagram of a processing system **1000**, according to an embodiment. In various embodiments, the system **1000** includes one or more processors **1002** and one or more graphics processors **1008**, and may be a single processor desktop system, a multiprocessor workstation system, or a server system having a large number of processors **1002** or processor cores **1007**. In one embodiment, the system **1000** is a processing platform incorporated within a system-on-a-chip (SoC) integrated circuit for use in mobile, handheld, or embedded devices.

An embodiment of system **1000** can include, or be incorporated within a server-based gaming platform, a game console, including a game and media console, a mobile gaming console, a handheld game console, or an online game console. In some embodiments system **1000** is a mobile phone, smart phone, tablet computing device or mobile Internet device. Data processing system **1000** can also include, couple with, or be integrated within a wearable device, such as a smart watch wearable device, smart eyewear device, augmented reality device, or virtual reality device. In some embodiments, data processing system **1000** is a television or set top box device having one or more processors **1002** and a graphical interface generated by one or more graphics processors **1008**.

In some embodiments, the one or more processors **1002** each include one or more processor cores **1007** to process instructions which, when executed, perform operations for system and user software. In some embodiments, each of the one or more processor cores **1007** is configured to process a specific instruction set **1009**. In some embodiments, instruction set **1009** may facilitate Complex Instruction Set Computing (CISC), Reduced Instruction Set Computing (RISC), or computing via a Very Long Instruction Word (VLIW). Multiple processor cores **1007** may each process a different instruction set **1009**, which may include instructions to facilitate the emulation of other instruction sets. Processor core **1007** may also include other processing devices, such as a Digital Signal Processor (DSP).

In some embodiments, the processor **1002** includes cache memory **1004**. Depending on the architecture, the processor **1002** can have a single internal cache or multiple levels of internal cache. In some embodiments, the cache memory is shared among various components of the processor **1002**. In some embodiments, the processor **1002** also uses an external cache (e.g., a Level-3 (L3) cache or Last Level Cache (LLC)) (not shown), which may be shared among processor cores **1007** using known cache coherency techniques. A register file **1006** is additionally included in processor **1002** which may include different types of registers for storing different types of data (e.g., integer registers, floating point registers, status registers, and an instruction pointer register). Some registers may be general-purpose registers, while other registers may be specific to the design of the processor **1002**.

In some embodiments, processor **1002** is coupled to a processor bus **1010** to transmit communication signals such as address, data, or control signals between processor **1002** and other components in system **1000**. In one embodiment, the system **1000** uses an exemplary ‘hub’ system architecture, including a memory controller hub **1016** and an Input Output (I/O) controller hub **1030**. A memory controller hub **1016** facilitates communication between a memory device

and other components of system 1000, while an I/O Controller Hub (ICH) 1030 provides connections to I/O devices via a local I/O bus. In one embodiment, the logic of the memory controller hub 1016 is integrated within the processor.

Memory device 1020 can be a dynamic random access memory (DRAM) device, a static random access memory (SRAM) device, flash memory device, phase-change memory device, or some other memory device having suitable performance to serve as process memory. In one embodiment, the memory device 1020 can operate as system memory for the system 1000, to store data 1022 and instructions 1021 for use when the one or more processors 1002 executes an application or process. Memory controller hub 1016 also couples with an optional external graphics processor 1012, which may communicate with the one or more graphics processors 1008 in processors 1002 to perform graphics and media operations.

In some embodiments, ICH 1030 enables peripherals to connect to memory device 1020 and processor 1002 via a high-speed I/O bus. The I/O peripherals include, but are not limited to, an audio controller 1046, a firmware interface 1028, a wireless transceiver 1026 (e.g., Wi-Fi, Bluetooth), a data storage device 1024 (e.g., hard disk drive, flash memory, etc.), and a legacy I/O controller 1040 for coupling legacy (e.g., Personal System 2 (PS/2)) devices to the system. One or more Universal Serial Bus (USB) controllers 1042 connect input devices, such as keyboard and mouse 1044 combinations. A network controller 1034 may also couple to ICH 1030. In some embodiments, a high-performance network controller (not shown) couples to processor bus 1010. It will be appreciated that the system 1000 shown is exemplary and not limiting, as other types of data processing systems that are differently configured may also be used. For example, the I/O controller hub 1030 may be integrated within the one or more processor 1002, or the memory controller hub 1016 and I/O controller hub 1030 may be integrated into a discreet external graphics processor, such as the external graphics processor 1012.

FIG. 11 is a block diagram of an embodiment of a processor 1100 having one or more processor cores 1102A-1102N, an integrated memory controller 1114, and an integrated graphics processor 1108. Those elements of FIG. 11 having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such. Processor 1100 can include additional cores up to and including additional core 1102N represented by the dashed lined boxes. Each of processor cores 1102A-1102N includes one or more internal cache units 1104A-1104N. In some embodiments, each processor core also has access to one or more shared cache units 1106.

The internal cache units 1104A-1104N and shared cache units 1106 represent a cache memory hierarchy within the processor 1100. The cache memory hierarchy may include at least one level of instruction and data cache within each processor core and one or more levels of shared mid-level cache, such as a Level 2 (L2), Level 3 (L3), Level 4 (L4), or other levels of cache, where the highest level of cache before external memory is classified as the LLC. In some embodiments, cache coherency logic maintains coherency between the various cache units 1106 and 1104A-1104N.

In some embodiments, processor 1100 may also include a set of one or more bus controller units 1116 and a system agent core 1110. The one or more bus controller units 1116 manage a set of peripheral buses, such as one or more

Peripheral Component Interconnect buses (e.g., PCI, PCI Express). System agent core 1110 provides management functionality for the various processor components. In some embodiments, system agent core 1110 includes one or more integrated memory controllers 1114 to manage access to various external memory devices (not shown).

In some embodiments, one or more of the processor cores 1102A-1102N include support for simultaneous multi-threading. In such embodiment, the system agent core 1110 includes components for coordinating and operating cores 1102A-1102N during multi-threaded processing. System agent core 1110 may additionally include a power control unit (PCU), which includes logic and components to regulate the power state of processor cores 1102A-1102N and graphics processor 1108.

In some embodiments, processor 1100 additionally includes graphics processor 1108 to execute graphics processing operations. In some embodiments, the graphics processor 1108 couples with the set of shared cache units 1106, and the system agent core 1110, including the one or more integrated memory controllers 1114. In some embodiments, a display controller 1111 is coupled with the graphics processor 1108 to drive graphics processor output to one or more coupled displays. In some embodiments, display controller 1111 may be a separate module coupled with the graphics processor via at least one interconnect, or may be integrated within the graphics processor 1108 or system agent core 1110.

In some embodiments, a ring based interconnect unit 1112 is used to couple the internal components of the processor 1100. However, an alternative interconnect unit may be used, such as a point-to-point interconnect, a switched interconnect, or other techniques, including techniques well known in the art. In some embodiments, graphics processor 1108 couples with the ring interconnect 1112 via an I/O link 1113.

The exemplary I/O link 1113 represents at least one of multiple varieties of I/O interconnects, including an on-package I/O interconnect which facilitates communication between various processor components and a high-performance embedded memory module 1118, such as an eDRAM module. In some embodiments, each of the processor cores 1102-1102N and graphics processor 1108 use embedded memory modules 1118 as a shared Last Level Cache.

In some embodiments, processor cores 1102A-1102N are homogenous cores executing the same instruction set architecture. In another embodiment, processor cores 1102A-1102N are heterogeneous in terms of instruction set architecture (ISA), where one or more of processor cores 1102A-N execute a first instruction set, while at least one of the other cores executes a subset of the first instruction set or a different instruction set. In one embodiment processor cores 1102A-1102N are heterogeneous in terms of micro-architecture, where one or more cores having a relatively higher power consumption couple with one or more power cores having a lower power consumption. Additionally, processor 1100 can be implemented on one or more chips or as an SoC integrated circuit having the illustrated components, in addition to other components.

FIG. 12 is a block diagram of a graphics processor 1200, which may be a discrete graphics processing unit, or may be a graphics processor integrated with a plurality of processing cores. In some embodiments, the graphics processor communicates via a memory mapped I/O interface to registers on the graphics processor and with commands placed into the processor memory. In some embodiments, graphics processor 1200 includes a memory interface 1214 to access

memory. Memory interface **1214** can be an interface to local memory, one or more internal caches, one or more shared external caches, and/or to system memory.

In some embodiments, graphics processor **1200** also includes a display controller **1202** to drive display output data to a display device **1220**. Display controller **1202** includes hardware for one or more overlay planes for the display and composition of multiple layers of video or user interface elements. In some embodiments, graphics processor **1200** includes a video codec engine **1206** to encode, decode, or transcode media to, from, or between one or more media encoding formats, including, but not limited to Moving Picture Experts Group (MPEG) formats such as MPEG-2, Advanced Video Coding (AVC) formats such as H.264/MPEG-4 AVC, as well as the Society of Motion Picture & Television Engineers (SMPTE) 421M/VC-1, and Joint Photographic Experts Group (JPEG) formats such as JPEG, and Motion JPEG (MJPEG) formats.

In some embodiments, graphics processor **1200** includes a block image transfer (BLIT) engine **1204** to perform two-dimensional (2D) rasterizer operations including, for example, bit-boundary block transfers. However, in one embodiment, 2D graphics operations are performed using one or more components of graphics processing engine (GPE) **1210**. In some embodiments, graphics processing engine **1210** is a compute engine for performing graphics operations, including three-dimensional (3D) graphics operations and media operations.

In some embodiments, GPE **1210** includes a 3D pipeline **1212** for performing 3D operations, such as rendering three-dimensional images and scenes using processing functions that act upon 3D primitive shapes (e.g., rectangle, triangle, etc.). The 3D pipeline **1212** includes programmable and fixed function elements that perform various tasks within the element and/or spawn execution threads to a 3D/Media sub-system **1215**. While 3D pipeline **1212** can be used to perform media operations, an embodiment of GPE **1210** also includes a media pipeline **1216** that is specifically used to perform media operations, such as video post-processing and image enhancement.

In some embodiments, media pipeline **1216** includes fixed function or programmable logic units to perform one or more specialized media operations, such as video decode acceleration, video de-interlacing, and video encode acceleration in place of, or on behalf of video codec engine **1206**. In some embodiments, media pipeline **1216** additionally includes a thread spawning unit to spawn threads for execution on 3D/Media sub-system **1215**. The spawned threads perform computations for the media operations on one or more graphics execution units included in 3D/Media sub-system **1215**.

In some embodiments, 3D/Media subsystem **1215** includes logic for executing threads spawned by 3D pipeline **1212** and media pipeline **1216**. In one embodiment, the pipelines send thread execution requests to 3D/Media subsystem **1215**, which includes thread dispatch logic for arbitrating and dispatching the various requests to available thread execution resources. The execution resources include an array of graphics execution units to process the 3D and media threads. In some embodiments, 3D/Media subsystem **1215** includes one or more internal caches for thread instructions and data. In some embodiments, the subsystem also includes shared memory, including registers and addressable memory, to share data between threads and to store output data.

3D/Media Processing

FIG. **13** is a block diagram of a graphics processing engine **1310** of a graphics processor in accordance with some embodiments. In one embodiment, the graphics processing engine (GPE) **1310** is a version of the GPE **1210** shown in FIG. **12**. Elements of FIG. **13** having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such. For example, the 3D pipeline **1212** and media pipeline **1216** of FIG. **12** are illustrated. The media pipeline **1216** is optional in some embodiments of the GPE **1310** and may not be explicitly included within the GPE **1310**. For example, and in at least one embodiment, a separate media and/or image processor is coupled to the GPE **1310**.

In some embodiments, GPE **1310** couples with or includes a command streamer **1303**, which provides a command stream to the 3D pipeline **1212** and/or media pipelines **1216**. In some embodiments, command streamer **1303** is coupled with memory, which can be system memory, or one or more of internal cache memory and shared cache memory. In some embodiments, command streamer **1303** receives commands from the memory and sends the commands to 3D pipeline **1212** and/or media pipeline **1216**. The commands are directives fetched from a ring buffer, which stores commands for the 3D pipeline **1212** and media pipeline **1216**. In one embodiment, the ring buffer can additionally include batch command buffers storing batches of multiple commands. The commands for the 3D pipeline **1212** can also include references to data stored in memory, such as but not limited to vertex and geometry data for the 3D pipeline **1212** and/or image data and memory objects for the media pipeline **1216**. The 3D pipeline **1212** and media pipeline **1216** process the commands and data by performing operations via logic within the respective pipelines or by dispatching one or more execution threads to a graphics core array **1314**.

In various embodiments, the 3D pipeline **1212** can execute one or more shader programs, such as vertex shaders, geometry shaders, pixel shaders, fragment shaders, compute shaders, or other shader programs, by processing the instructions and dispatching execution threads to the graphics core array **1314**. The graphics core array **1314** provides a unified block of execution resources. Multi-purpose execution logic (e.g., execution units) within the graphic core array **1314** includes support for various 3D API shader languages and can execute multiple simultaneous execution threads associated with multiple shaders.

In some embodiments, the graphics core array **1314** also includes execution logic to perform media functions, such as video and/or image processing. In one embodiment, the execution units additionally include general-purpose logic that is programmable to perform parallel general purpose computational operations, in addition to graphics processing operations. The general-purpose logic can perform processing operations in parallel or in conjunction with general purpose logic within the processor core(s) **1007** of FIG. **10** or core **1102A-1102N** as in FIG. **11**.

Output data generated by threads executing on the graphics core array **1314** can output data to memory in a unified return buffer (URB) **1318**. The URB **1318** can store data for multiple threads. In some embodiments, the URB **1318** may be used to send data between different threads executing on the graphics core array **1314**. In some embodiments, the URB **1318** may additionally be used for synchronization between threads on the graphics core array and fixed function logic within the shared function logic **1320**.

In some embodiments, graphics core array **1314** is scalable, such that the array includes a variable number of graphics cores, each having a variable number of execution units based on the target power and performance level of GPE **1310**. In one embodiment, the execution resources are dynamically scalable, such that execution resources may be enabled or disabled as needed.

The graphics core array **1314** couples with shared function logic **1320** that includes multiple resources that are shared between the graphics cores in the graphics core array. The shared functions within the shared function logic **1320** are hardware logic units that provide specialized supplemental functionality to the graphics core array **1314**. In various embodiments, shared function logic **1320** includes but is not limited to sampler **1321**, math **1322**, and inter-thread communication (ITC) **1323** logic. Additionally, some embodiments implement one or more cache(s) **1325** within the shared function logic **1320**. A shared function is implemented where the demand for a given specialized function is insufficient for inclusion within the graphics core array **1314**. Instead a single instantiation of that specialized function is implemented as a stand-alone entity in the shared function logic **1320** and shared among the execution resources within the graphics core array **1314**. The precise set of functions that are shared between the graphics core array **1314** and included within the graphics core array **1314** varies between embodiments.

FIG. **14** is a block diagram of another embodiment of a graphics processor **1400**. Elements of FIG. **14** having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such.

In some embodiments, graphics processor **1400** includes a ring interconnect **1402**, a pipeline front-end **1404**, a media engine **1437**, and graphics cores **1480A-1480N**. In some embodiments, ring interconnect **1402** couples the graphics processor to other processing units, including other graphics processors or one or more general-purpose processor cores. In some embodiments, the graphics processor is one of many processors integrated within a multi-core processing system.

In some embodiments, graphics processor **1400** receives batches of commands via ring interconnect **1402**. The incoming commands are interpreted by a command streamer **1403** in the pipeline front-end **1404**. In some embodiments, graphics processor **1400** includes scalable execution logic to perform 3D geometry processing and media processing via the graphics core(s) **1480A-1480N**. For 3D geometry processing commands, command streamer **1403** supplies commands to geometry pipeline **1436**. For at least some media processing commands, command streamer **1403** supplies the commands to a video front end **1434**, which couples with a media engine **1437**. In some embodiments, media engine **1437** includes a Video Quality Engine (VQE) **1430** for video and image post-processing and a multi-format encode/decode (MFX) **1433** engine to provide hardware-accelerated media data encode and decode. In some embodiments, geometry pipeline **1436** and media engine **1437** each generate execution threads for the thread execution resources provided by at least one graphics core **1480A**.

In some embodiments, graphics processor **1400** includes scalable thread execution resources featuring modular cores **1480A-1480N** (sometimes referred to as core slices), each having multiple sub-cores **1450A-1450N**, **1460A-1460N** (sometimes referred to as core sub-slices). In some embodiments, graphics processor **1400** can have any number of graphics cores **1480A** through **1480N**. In some embodi-

ments, graphics processor **1400** includes a graphics core **1480A** having at least a first sub-core **1450A** and a second core sub-core **1460A**. In other embodiments, the graphics processor is a low power processor with a single sub-core (e.g., **1450A**). In some embodiments, graphics processor **1400** includes multiple graphics cores **1480A-1480N**, each including a set of first sub-cores **1450A-1450N** and a set of second sub-cores **1460A-1460N**. Each sub-core in the set of first sub-cores **1450A-1450N** includes at least a first set of execution units **1452A-1452N** and media/texture samplers **1454A-1454N**. Each sub-core in the set of second sub-cores **1460A-1460N** includes at least a second set of execution units **1462A-1462N** and samplers **1464A-1464N**. In some embodiments, each sub-core **1450A-1450N**, **1460A-1460N** shares a set of shared resources **1470A-1470N**. In some embodiments, the shared resources include shared cache memory and pixel operation logic. Other shared resources may also be included in the various embodiments of the graphics processor.

Execution Logic

FIG. **15** illustrates thread execution logic **1500** including an array of processing elements employed in some embodiments of a GPE. Elements of FIG. **15** having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such.

In some embodiments, thread execution logic **1500** includes a pixel shader **1502**, a thread dispatcher **1504**, instruction cache **1506**, a scalable execution unit array including a plurality of execution units **1508A-608N**, a sampler **1510**, a data cache **1512**, and a data port **1514**. In one embodiment, the included components are interconnected via an interconnect fabric that links to each of the components. In some embodiments, thread execution logic **1500** includes one or more connections to memory, such as system memory or cache memory, through one or more of instruction cache **1506**, data port **1514**, sampler **1510**, and execution unit array **1508A-1508N**. In some embodiments, each execution unit (e.g. **1508A**) is an individual vector processor capable of executing multiple simultaneous threads and processing multiple data elements in parallel for each thread. In some embodiments, execution unit array **1508A-1508N** includes any number individual execution units.

In some embodiments, execution unit array **1508A-1508N** is primarily used to execute “shader” programs. In some embodiments, the execution units in array **1508A-1508N** execute an instruction set that includes native support for many standard 3D graphics shader instructions, such that shader programs from graphics libraries (e.g., Direct 3D and OpenGL) are executed with a minimal translation. The execution units support vertex and geometry processing (e.g., vertex programs, geometry programs, vertex shaders), pixel processing (e.g., pixel shaders, fragment shaders) and general-purpose processing (e.g., compute and media shaders).

Each execution unit in execution unit array **1508A-1508N** operates on arrays of data elements. The number of data elements is the “execution size,” or the number of channels for the instruction. An execution channel is a logical unit of execution for data element access, masking, and flow control within instructions. The number of channels may be independent of the number of physical Arithmetic Logic Units (ALUs) or Floating Point Units (FPUs) for a particular graphics processor. In some embodiments, execution units **1508A-1508N** support integer and floating-point data types.

The execution unit instruction set includes single instruction multiple data (SIMD) or single instruction multiple thread (SIMT) instructions. The various data elements can be stored as a packed data type in a register and the execution unit will process the various elements based on the data size of the elements. For example, when operating on a 256-bit wide vector, the 256 bits of the vector are stored in a register and the execution unit operates on the vector as four separate 64-bit packed data elements (Quad-Word (QW) size data elements), eight separate 32-bit packed data elements (Double Word (DW) size data elements), sixteen separate 16-bit packed data elements (Word (W) size data elements), or thirty-two separate 8-bit data elements (byte (B) size data elements). However, different vector widths and register sizes are possible.

One or more internal instruction caches (e.g., **1506**) are included in the thread execution logic **1500** to cache thread instructions for the execution units. In some embodiments, one or more data caches (e.g., **1512**) are included to cache thread data during thread execution. In some embodiments, sampler **1510** is included to provide texture sampling for 3D operations and media sampling for media operations. In some embodiments, sampler **1510** includes specialized texture or media sampling functionality to process texture or media data during the sampling process before providing the sampled data to an execution unit.

During execution, the graphics and media pipelines send thread initiation requests to thread execution logic **1500** via thread spawning and dispatch logic. In some embodiments, thread execution logic **1500** includes a local thread dispatcher **1504** that arbitrates thread initiation requests from the graphics and media pipelines and instantiates the requested threads on one or more execution units **1508A-1508N**. For example, the geometry pipeline (e.g., **1436** of FIG. **14**) dispatches vertex processing, tessellation, or geometry processing threads to thread execution logic **1500** (FIG. **15**). In some embodiments, thread dispatcher **1504** can also process runtime thread spawning requests from the executing shader programs.

Once a group of geometric objects has been processed and rasterized into pixel data, pixel shader **1502** is invoked to further compute output information and cause results to be written to output surfaces (e.g., color buffers, depth buffers, stencil buffers, etc.). In some embodiments, pixel shader **1502** calculates the values of the various vertex attributes that are to be interpolated across the rasterized object. In some embodiments, pixel shader **1502** then executes an application programming interface (API)-supplied pixel shader program. To execute the pixel shader program, pixel shader **1502** dispatches threads to an execution unit (e.g., **1508A**) via thread dispatcher **1504**. In some embodiments, pixel shader **1502** uses texture sampling logic in sampler **1510** to access texture data in texture maps stored in memory. Arithmetic operations on the texture data and the input geometry data compute pixel color data for each geometric fragment, or discards one or more pixels from further processing.

In some embodiments, the data port **1514** provides a memory access mechanism for the thread execution logic **1500** output processed data to memory for processing on a graphics processor output pipeline. In some embodiments, the data port **1514** includes or couples to one or more cache memories (e.g., data cache **1512**) to cache data for memory access via the data port.

FIG. **16** is a block diagram illustrating a graphics processor instruction formats **1600** according to some embodiments. In one or more embodiment, the graphics processor

execution units support an instruction set having instructions in multiple formats. The solid lined boxes illustrate the components that are generally included in an execution unit instruction, while the dashed lines include components that are optional or that are only included in a sub-set of the instructions. In some embodiments, instruction format **1600** described and illustrated are macro-instructions, in that they are instructions supplied to the execution unit, as opposed to micro-operations resulting from instruction decode once the instruction is processed.

In some embodiments, the graphics processor execution units natively support instructions in a 128-bit instruction format **1610**. A 64-bit compacted instruction format **1630** is available for some instructions based on the selected instruction, instruction options, and number of operands. The native 128-bit instruction format **1610** provides access to all instruction options, while some options and operations are restricted in the 64-bit instruction format **1630**. The native instructions available in the 64-bit instruction format **1630** vary by embodiment. In some embodiments, the instruction is compacted in part using a set of index values in an index field **1613**. The execution unit hardware references a set of compaction tables based on the index values and uses the compaction table outputs to reconstruct a native instruction in the 128-bit instruction format **1610**.

For each format, instruction opcode **1612** defines the operation that the execution unit is to perform. The execution units execute each instruction in parallel across the multiple data elements of each operand. For example, in response to an add instruction the execution unit performs a simultaneous add operation across each color channel representing a texture element or picture element. By default, the execution unit performs each instruction across all data channels of the operands. In some embodiments, instruction control field **1614** enables control over certain execution options, such as channels selection (e.g., predication) and data channel order (e.g., swizzle). For 128-bit instructions **1610** an exec-size field **1616** limits the number of data channels that will be executed in parallel. In some embodiments, exec-size field **1616** is not available for use in the 64-bit compact instruction format **1630**.

Some execution unit instructions have up to three operands including two source operands, src0 **1620**, src1 **1622**, and one destination **1618**. In some embodiments, the execution units support dual destination instructions, where one of the destinations is implied. Data manipulation instructions can have a third source operand (e.g., SRC2 **1624**), where the instruction opcode **1612** determines the number of source operands. An instruction's last source operand can be an immediate (e.g., hard-coded) value passed with the instruction.

In some embodiments, the 128-bit instruction format **1610** includes an access/address mode information **1626** specifying, for example, whether direct register addressing mode or indirect register addressing mode is used. When direct register addressing mode is used, the register address of one or more operands is directly provided by bits in the instruction **1610**.

In some embodiments, the 128-bit instruction format **1610** includes an access/address mode field **1626**, which specifies an address mode and/or an access mode for the instruction. In one embodiment, the access mode to define a data access alignment for the instruction. Some embodiments support access modes including a 16-byte aligned access mode and a 1-byte aligned access mode, where the byte alignment of the access mode determines the access alignment of the instruction operands. For example, when in

a first mode, the instruction **1610** may use byte-aligned addressing for source and destination operands and when in a second mode, the instruction **1610** may use 16-byte-aligned addressing for all source and destination operands.

In one embodiment, the address mode portion of the access/address mode field **1626** determines whether the instruction is to use direct or indirect addressing. When direct register addressing mode is used bits in the instruction **1610** directly provide the register address of one or more operands. When indirect register addressing mode is used, the register address of one or more operands may be computed based on an address register value and an address immediate field in the instruction.

In some embodiments, instructions are grouped based on opcode **1612** bit-fields to simplify Opcode decode **1640**. For an 8-bit opcode, bits **4**, **5**, and **6** allow the execution unit to determine the type of opcode. The precise opcode grouping shown is merely an example. In some embodiments, a move and logic opcode group **1642** includes data movement and logic instructions (e.g., move (mov), compare (cmp)). In some embodiments, move and logic group **1642** shares the five most significant bits (MSB), where move (mov) instructions are in the form of 0000xxxxb and logic instructions are in the form of 0001xxxxb. A flow control instruction group **1644** (e.g., call, jump (jmp)) includes instructions in the form of 0010xxxxb (e.g., 0x20). A miscellaneous instruction group **1646** includes a mix of instructions, including synchronization instructions (e.g., wait, send) in the form of 0011xxxxb (e.g., 0x30). A parallel math instruction group **1648** includes component-wise arithmetic instructions (e.g., add, multiply (mul)) in the form of 0100xxxxb (e.g., 0x40). The parallel math group **1648** performs the arithmetic operations in parallel across data channels. The vector math group **1650** includes arithmetic instructions (e.g., dp4) in the form of 0101xxxxb (e.g., 0x50). The vector math group performs arithmetic such as dot product calculations on vector operands.

Graphics Pipeline

FIG. **17** is a block diagram of another embodiment of a graphics processor **1700**. Elements of FIG. **17** having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such.

In some embodiments, graphics processor **1700** includes a graphics pipeline **1720**, a media pipeline **1730**, a display engine **1740**, thread execution logic **1750**, and a render output pipeline **1770**. In some embodiments, graphics processor **1700** is a graphics processor within a multi-core processing system that includes one or more general-purpose processing cores. The graphics processor is controlled by register writes to one or more control registers (not shown) or via commands issued to graphics processor **1700** via a ring interconnect **1702**. In some embodiments, ring interconnect **1702** couples graphics processor **1700** to other processing components, such as other graphics processors or general-purpose processors. Commands from ring interconnect **1702** are interpreted by a command streamer **1703**, which supplies instructions to individual components of graphics pipeline **1720** or media pipeline **1730**.

In some embodiments, command streamer **1703** directs the operation of a vertex fetcher **1705** that reads vertex data from memory and executes vertex-processing commands provided by command streamer **1703**. In some embodiments, vertex fetcher **1705** provides vertex data to a vertex shader **1707**, which performs coordinate space transformation and lighting operations to each vertex. In some embodi-

ments, vertex fetcher **1705** and vertex shader **1707** execute vertex-processing instructions by dispatching execution threads to execution units **1752A**, **1752B** via a thread dispatcher **1731**.

In some embodiments, execution units **1752A**, **1752B** are an array of vector processors having an instruction set for performing graphics and media operations. In some embodiments, execution units **1752A**, **1752B** have an attached L1 cache **1751** that is specific for each array or shared between the arrays. The cache can be configured as a data cache, an instruction cache, or a single cache that is partitioned to contain data and instructions in different partitions.

In some embodiments, graphics pipeline **1720** includes tessellation components to perform hardware-accelerated tessellation of 3D objects. In some embodiments, a programmable hull shader **1711** configures the tessellation operations. A programmable domain shader **1717** provides back-end evaluation of tessellation output. A tessellator **1713** operates at the direction of hull shader **1711** and contains special purpose logic to generate a set of detailed geometric objects based on a coarse geometric model that is provided as input to graphics pipeline **1720**. In some embodiments, if tessellation is not used, tessellation components **1711**, **1713**, **1717** can be bypassed.

In some embodiments, complete geometric objects can be processed by a geometry shader **1719** via one or more threads dispatched to execution units **1752A**, **1752B**, or can proceed directly to the clipper **1729**. In some embodiments, the geometry shader operates on entire geometric objects, rather than vertices or patches of vertices as in previous stages of the graphics pipeline. If the tessellation is disabled the geometry shader **1719** receives input from the vertex shader **1707**. In some embodiments, geometry shader **1719** is programmable by a geometry shader program to perform geometry tessellation if the tessellation units are disabled.

Before rasterization, a clipper **1729** processes vertex data. The clipper **1729** may be a fixed function clipper or a programmable clipper having clipping and geometry shader functions. In some embodiments, a rasterizer and depth test component **1773** in the render output pipeline **1770** dispatches pixel shaders to convert the geometric objects into their per pixel representations. In some embodiments, pixel shader logic is included in thread execution logic **1750**. In some embodiments, an application can bypass rasterization and access un-rasterized vertex data via a stream out unit **1723**.

The graphics processor **1700** has an interconnect bus, interconnect fabric, or some other interconnect mechanism that allows data and message passing amongst the major components of the processor. In some embodiments, execution units **1752A**, **1752B** and associated cache(s) **1751**, texture and media sampler **1754**, and texture/sampler cache **1758** interconnect via a data port **1756** to perform memory access and communicate with render output pipeline components of the processor. In some embodiments, sampler **1754**, caches **1751**, **1758** and execution units **1752A**, **1752B** each have separate memory access paths.

In some embodiments, render output pipeline **1770** contains a rasterizer and depth test component **1773** that converts vertex-based objects into an associated pixel-based representation. In some embodiments, the render output pipeline **1770** includes a windower/masker unit to perform fixed function triangle and line rasterization. An associated render cache **1778** and depth cache **1779** are also available in some embodiments. A pixel operations component **1777** performs pixel-based operations on the data, though in some instances, pixel operations associated with 2D operations

(e.g. bit block image transfers with blending) are performed by the 2D engine **1741**, or substituted at display time by the display controller **1743** using overlay display planes. In some embodiments, a shared L3 cache **1775** is available to all graphics components, allowing the sharing of data without the use of main system memory.

In some embodiments, graphics processor media pipeline **1730** includes a media engine **1737** and a video front end **1734**. In some embodiments, video front end **1734** receives pipeline commands from the command streamer **1703**. In some embodiments, media pipeline **1730** includes a separate command streamer. In some embodiments, video front-end **1734** processes media commands before sending the command to the media engine **1737**. In some embodiments, media engine **1737** includes thread spawning functionality to spawn threads for dispatch to thread execution logic **1750** via thread dispatcher **1731**.

In some embodiments, graphics processor **1700** includes a display engine **1740**. In some embodiments, display engine **1740** is external to processor **1700** and couples with the graphics processor via the ring interconnect **1702**, or some other interconnect bus or fabric. In some embodiments, display engine **1740** includes a 2D engine **1741** and a display controller **1743**. In some embodiments, display engine **1740** contains special purpose logic capable of operating independently of the 3D pipeline. In some embodiments, display controller **1743** couples with a display device (not shown), which may be a system integrated display device, as in a laptop computer, or an external display device attached via a display device connector.

In some embodiments, graphics pipeline **1720** and media pipeline **1730** are configurable to perform operations based on multiple graphics and media programming interfaces and are not specific to any one application programming interface (API). In some embodiments, driver software for the graphics processor translates API calls that are specific to a particular graphics or media library into commands that can be processed by the graphics processor. In some embodiments, support is provided for the Open Graphics Library (OpenGL) and Open Computing Language (OpenCL) from the Khronos Group, the Direct3D library from the Microsoft Corporation, or support may be provided to both OpenGL and D3D. Support may also be provided for the Open Source Computer Vision Library (OpenCV). A future API with a compatible 3D pipeline would also be supported if a mapping can be made from the pipeline of the future API to the pipeline of the graphics processor.

Graphics Pipeline Programming

FIG. **18A** is a block diagram illustrating a graphics processor command format **1800** according to some embodiments. FIG. **18B** is a block diagram illustrating a graphics processor command sequence **1810** according to an embodiment. The solid lined boxes in FIG. **18A** illustrate the components that are generally included in a graphics command while the dashed lines include components that are optional or that are only included in a sub-set of the graphics commands. The exemplary graphics processor command format **1800** of FIG. **18A** includes data fields to identify a target client **1802** of the command, a command operation code (opcode) **1804**, and the relevant data **1806** for the command. A sub-opcode **1805** and a command size **1808** are also included in some commands.

In some embodiments, client **1802** specifies the client unit of the graphics device that processes the command data. In some embodiments, a graphics processor command parser examines the client field of each command to condition the further processing of the command and route the command

data to the appropriate client unit. In some embodiments, the graphics processor client units include a memory interface unit, a render unit, a 2D unit, a 3D unit, and a media unit. Each client unit has a corresponding processing pipeline that processes the commands. Once the command is received by the client unit, the client unit reads the opcode **1804** and, if present, sub-opcode **1805** to determine the operation to perform. The client unit performs the command using information in data field **1806**. For some commands an explicit command size **1808** is expected to specify the size of the command. In some embodiments, the command parser automatically determines the size of at least some of the commands based on the command opcode. In some embodiments, commands are aligned via multiples of a double word.

The flow diagram in FIG. **18B** shows an exemplary graphics processor command sequence **1810**. In some embodiments, software or firmware of a data processing system that features an embodiment of a graphics processor uses a version of the command sequence shown to set up, execute, and terminate a set of graphics operations. A sample command sequence is shown and described for purposes of example only as embodiments are not limited to these specific commands or to this command sequence. Moreover, the commands may be issued as batch of commands in a command sequence, such that the graphics processor will process the sequence of commands in at least partially concurrence.

In some embodiments, the graphics processor command sequence **1810** may begin with a pipeline flush command **1812** to cause any active graphics pipeline to complete the currently pending commands for the pipeline. In some embodiments, the 3D pipeline **1822** and the media pipeline **1824** do not operate concurrently. The pipeline flush is performed to cause the active graphics pipeline to complete any pending commands. In response to a pipeline flush, the command parser for the graphics processor will pause command processing until the active drawing engines complete pending operations and the relevant read caches are invalidated. Optionally, any data in the render cache that is marked 'dirty' can be flushed to memory. In some embodiments, pipeline flush command **1812** can be used for pipeline synchronization or before placing the graphics processor into a low power state.

In some embodiments, a pipeline select command **1813** is used when a command sequence requires the graphics processor to explicitly switch between pipelines. In some embodiments, a pipeline select command **1813** is required only once within an execution context before issuing pipeline commands unless the context is to issue commands for both pipelines. In some embodiments, a pipeline flush command **1812** is required immediately before a pipeline switch via the pipeline select command **1813**.

In some embodiments, a pipeline control command **1814** configures a graphics pipeline for operation and is used to program the 3D pipeline **1822** and the media pipeline **1824**. In some embodiments, pipeline control command **1814** configures the pipeline state for the active pipeline. In one embodiment, the pipeline control command **1814** is used for pipeline synchronization and to clear data from one or more cache memories within the active pipeline before processing a batch of commands.

In some embodiments, commands for the return buffer state **1816** are used to configure a set of return buffers for the respective pipelines to write data. Some pipeline operations require the allocation, selection, or configuration of one or more return buffers into which the operations write inter-

mediate data during processing. In some embodiments, the graphics processor also uses one or more return buffers to store output data and to perform cross thread communication. In some embodiments, configuring the return buffer state **1816** includes selecting the size and number of return buffers to use for a set of pipeline operations.

The remaining commands in the command sequence differ based on the active pipeline for operations. Based on a pipeline determination **1820**, the command sequence is tailored to the 3D pipeline **1822** beginning with the 3D pipeline state **1830**, or the media pipeline **1824** beginning at the media pipeline state **1840**.

The commands for the 3D pipeline state **1830** include 3D state setting commands for vertex buffer state, vertex element state, constant color state, depth buffer state, and other state variables that are to be configured before 3D primitive commands are processed. The values of these commands are determined at least in part based the particular 3D API in use. In some embodiments, 3D pipeline state **1830** commands are also able to selectively disable or bypass certain pipeline elements if those elements will not be used.

In some embodiments, 3D primitive **1832** command is used to submit 3D primitives to be processed by the 3D pipeline. Commands and associated parameters that are passed to the graphics processor via the 3D primitive **1832** command are forwarded to the vertex fetch function in the graphics pipeline. The vertex fetch function uses the 3D primitive **1832** command data to generate vertex data structures. The vertex data structures are stored in one or more return buffers. In some embodiments, 3D primitive **1832** command is used to perform vertex operations on 3D primitives via vertex shaders. To process vertex shaders, 3D pipeline **1822** dispatches shader execution threads to graphics processor execution units.

In some embodiments, 3D pipeline **1822** is triggered via an execute **1834** command or event. In some embodiments, a register write triggers command execution. In some embodiments execution is triggered via a ‘go’ or ‘kick’ command in the command sequence. In one embodiment command execution is triggered using a pipeline synchronization command to flush the command sequence through the graphics pipeline. The 3D pipeline will perform geometry processing for the 3D primitives. Once operations are complete, the resulting geometric objects are rasterized and the pixel engine colors the resulting pixels. Additional commands to control pixel shading and pixel back end operations may also be included for those operations.

In some embodiments, the graphics processor command sequence **1810** follows the media pipeline **1824** path when performing media operations. In general, the specific use and manner of programming for the media pipeline **1824** depends on the media or compute operations to be performed. Specific media decode operations may be offloaded to the media pipeline during media decode. In some embodiments, the media pipeline can also be bypassed and media decode can be performed in whole or in part using resources provided by one or more general-purpose processing cores. In one embodiment, the media pipeline also includes elements for general-purpose graphics processor unit (GPGPU) operations, where the graphics processor is used to perform SIMD vector operations using computational shader programs that are not explicitly related to the rendering of graphics primitives.

In some embodiments, media pipeline **1824** is configured in a similar manner as the 3D pipeline **1822**. A set of commands to configure the media pipeline state **1840** are dispatched or placed into a command queue before the

media object commands **1842**. In some embodiments, commands for the media pipeline state **1840** include data to configure the media pipeline elements that will be used to process the media objects. This includes data to configure the video decode and video encode logic within the media pipeline, such as encode or decode format. In some embodiments, commands for the media pipeline state **1840** also support the use of one or more pointers to “indirect” state elements that contain a batch of state settings.

In some embodiments, media object commands **1842** supply pointers to media objects for processing by the media pipeline. The media objects include memory buffers containing video data to be processed. In some embodiments, all media pipeline states must be valid before issuing a media object command **1842**. Once the pipeline state is configured and media object commands **1842** are queued, the media pipeline **1824** is triggered via an execute command **1844** or an equivalent execute event (e.g., register write). Output from media pipeline **1824** may then be post processed by operations provided by the 3D pipeline **1822** or the media pipeline **1824**. In some embodiments, GPGPU operations are configured and executed in a similar manner as media operations.

Graphics Software Architecture

FIG. **19** illustrates exemplary graphics software architecture for a data processing system **1900** according to some embodiments. In some embodiments, software architecture includes a 3D graphics application **1910**, an operating system **1920**, and at least one processor **1930**. In some embodiments, processor **1930** includes a graphics processor **1932** and one or more general-purpose processor core(s) **1934**. The graphics application **1910** and operating system **1920** each execute in the system memory **1950** of the data processing system.

In some embodiments, 3D graphics application **1910** contains one or more shader programs including shader instructions **1912**. The shader language instructions may be in a high-level shader language, such as the High Level Shader Language (HLSL) or the OpenGL Shader Language (GLSL). The application also includes executable instructions **1914** in a machine language suitable for execution by the general-purpose processor core(s) **1934**. The application also includes graphics objects **1916** defined by vertex data.

In some embodiments, operating system **1920** is a Microsoft® Windows® operating system from the Microsoft Corporation, a proprietary UNIX-like operating system, or an open source UNIX-like operating system using a variant of the Linux kernel. The operating system **1920** can support a graphics API **1922** such as the Direct3D API or the OpenGL API. When the Direct3D API is in use, the operating system **1920** uses a front-end shader compiler **1924** to compile any shader instructions **1912** in HLSL into a lower-level shader language. The compilation may be a just-in-time (JIT) compilation or the application can perform shader pre-compilation. In some embodiments, high-level shaders are compiled into low-level shaders during the compilation of the 3D graphics application **1910**.

In some embodiments, user mode graphics driver **1926** contains a back-end shader compiler **1927** to convert the shader instructions **1912** into a hardware specific representation. When the OpenGL API is in use, shader instructions **1912** in the GLSL high-level language are passed to a user mode graphics driver **1926** for compilation. In some embodiments, user mode graphics driver **1926** uses operating system kernel mode functions **1928** to communicate with a kernel mode graphics driver **1929**. In some embodi-

ments, kernel mode graphics driver **1929** communicates with graphics processor **1932** to dispatch commands and instructions.

IP Core Implementations

One or more aspects of at least one embodiment may be implemented by representative code stored on a machine-readable medium which represents and/or defines logic within an integrated circuit such as a processor. For example, the machine-readable medium may include instructions which represent various logic within the processor. When read by a machine, the instructions may cause the machine to fabricate the logic to perform the techniques described herein. Such representations, known as “IP cores,” are reusable units of logic for an integrated circuit that may be stored on a tangible, machine-readable medium as a hardware model that describes the structure of the integrated circuit. The hardware model may be supplied to various customers or manufacturing facilities, which load the hardware model on fabrication machines that manufacture the integrated circuit. The integrated circuit may be fabricated such that the circuit performs operations described in association with any of the embodiments described herein.

FIG. **20** is a block diagram illustrating an IP core development system **2000** that may be used to manufacture an integrated circuit to perform operations according to an embodiment. The IP core development system **2000** may be used to generate modular, re-usable designs that can be incorporated into a larger design or used to construct an entire integrated circuit (e.g., an SOC integrated circuit). A design facility **2030** can generate a software simulation **2010** of an IP core design in a high-level programming language (e.g., C/C++). The software simulation **2010** can be used to design, test, and verify the behavior of the IP core using a simulation model **2012**. The simulation model **2012** may include functional, behavioral, and/or timing simulations. A register transfer level (RTL) design **2015** can then be created or synthesized from the simulation model **2012**. The RTL design **2015** is an abstraction of the behavior of the integrated circuit that models the flow of digital signals between hardware registers, including the associated logic performed using the modeled digital signals. In addition to an RTL design **2015**, lower-level designs at the logic level or transistor level may also be created, designed, or synthesized. Thus, the particular details of the initial design and simulation may vary.

The RTL design **2015** or equivalent may be further synthesized by the design facility into a hardware model **2020**, which may be in a hardware description language (HDL), or some other representation of physical design data. The HDL may be further simulated or tested to verify the IP core design. The IP core design can be stored for delivery to a 3rd party fabrication facility **2065** using non-volatile memory **2040** (e.g., hard disk, flash memory, or any non-volatile storage medium). Alternatively, the IP core design may be transmitted (e.g., via the Internet) over a wired connection **2050** or wireless connection **2060**. The fabrication facility **2065** may then fabricate an integrated circuit that is based at least in part on the IP core design. The fabricated integrated circuit can be configured to perform operations in accordance with at least one embodiment described herein.

Exemplary System on a Chip Integrated Circuit

FIGS. **21-23** illustrate exemplary integrated circuits and associated graphics processors that may be fabricated using one or more IP cores, according to various embodiments described herein. In addition to what is illustrated, other logic and circuits may be included, including additional

graphics processors/cores, peripheral interface controllers, or general purpose processor cores.

FIG. **21** is a block diagram illustrating an exemplary system on a chip integrated circuit **2100** that may be fabricated using one or more IP cores, according to an embodiment. Exemplary integrated circuit **2100** includes one or more application processor(s) **2105** (e.g., CPUs), at least one graphics processor **2110**, and may additionally include an image processor **2115** and/or a video processor **2120**, any of which may be a modular IP core from the same or multiple different design facilities. Integrated circuit **2100** includes peripheral or bus logic including a USB controller **2125**, UART controller **2130**, an SPI/SDIO controller **2135**, and an I²S/I²C controller **2140**. Additionally, the integrated circuit can include a display device **2145** coupled to one or more of a high-definition multimedia interface (HDMI) controller **2150** and a mobile industry processor interface (MIPI) display interface **2155**. Storage may be provided by a flash memory subsystem **2160** including flash memory and a flash memory controller. Memory interface may be provided via a memory controller **2165** for access to SDRAM or SRAM memory devices. Some integrated circuits additionally include an embedded security engine **2170**.

FIG. **22** is a block diagram illustrating an exemplary graphics processor **2210** of a system on a chip integrated circuit that may be fabricated using one or more IP cores, according to an embodiment. Graphics processor **2210** can be a variant of the graphics processor **2110** of FIG. **21**. Graphics processor **2210** includes a vertex processor **2205** and one or more fragment processor(s) **2215A-2215N** (e.g., **2215A**, **2215B**, **2215C**, **2215D**, through **2215N-1**, and **2215N**). Graphics processor **2210** can execute different shader programs via separate logic, such that the vertex processor **2205** is optimized to execute operations for vertex shader programs, while the one or more fragment processor(s) **2215A-2215N** execute fragment (e.g., pixel) shading operations for fragment or pixel shader programs. The vertex processor **2205** performs the vertex processing stage of the 3D graphics pipeline and generates primitives and vertex data. The fragment processor(s) **2215A-2215N** use the primitive and vertex data generated by the vertex processor **2205** to produce a framebuffer that is displayed on a display device. In one embodiment, the fragment processor(s) **2215A-2215N** are optimized to execute fragment shader programs as provided for in the OpenGL API, which may be used to perform similar operations as a pixel shader program as provided for in the Direct 3D API.

Graphics processor **2210** additionally includes one or more memory management units (MMUs) **2220A-2220B**, cache(s) **2225A-2225B**, and circuit interconnect(s) **2230A-2230B**. The one or more MMU(s) **2220A-2220B** provide for virtual to physical address mapping for graphics processor **2210**, including for the vertex processor **2205** and/or fragment processor(s) **2215A-2215N**, which may reference vertex or image/texture data stored in memory, in addition to vertex or image/texture data stored in the one or more cache(s) **2225A-2225B**. In one embodiment, the one or more MMU(s) **2220A-2220B** may be synchronized with other MMUs within the system, including one or more MMUs associated with the one or more application processor(s) **2105**, image processor **2115**, and/or video processor **2120** of FIG. **21**, such that each processor **2105-2120** can participate in a shared or unified virtual memory system. The one or more circuit interconnect(s) **2230A-2230B** enable graphics processor **2210** to interface with other IP cores within the SoC, either via an internal bus of the SoC or via a direct connection, according to embodiments.

FIG. 23 is a block diagram illustrating an additional exemplary graphics processor 2310 of a system on a chip integrated circuit that may be fabricated using one or more IP cores, according to an embodiment. Graphics processor 2310 can be a variant of the graphics processor 2110 of FIG. 21. Graphics processor 2310 includes the one or more MMU(s) 2220A-2220B, cache(s) 2225A-2225B, and circuit interconnect(s) 2230A-2230B of the integrated circuit 2200 of FIG. 22.

Graphics processor 2310 includes one or more shader core(s) 2315A-2315N (e.g., 2315A, 2315B, 2315C, 2315D, 2315E, 2315F, through 2215N-1, and 2215N), which provides for a unified shader core architecture in which a single core or type or core can execute all types of programmable shader code, including shader program code to implement vertex shaders, fragment shaders, and/or compute shaders. The exact number of shader cores present can vary among embodiments and implementations. Additionally, graphics processor 2310 includes an inter-core task manager 2305, which acts as a thread dispatcher to dispatch execution threads to one or more shader core(s) 2315A-2315N. Graphics processor 2310 additionally includes a tiling unit 2318 to accelerate tiling operations for tile-based rendering, in which rendering operations for a scene are subdivided in image space. Tile-based rendering can be used to exploit local spatial coherence within a scene or to optimize use of internal caches.

References to “one embodiment”, “an embodiment”, “example embodiment”, “various embodiments”, etc., indicate that the embodiment(s) so described may include particular features, structures, or characteristics, but not every embodiment necessarily includes the particular features, structures, or characteristics. Further, some embodiments may have some, all, or none of the features described for other embodiments.

In the foregoing specification, embodiments have been described with reference to specific exemplary embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of embodiments as set forth in the appended claims. The Specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

In the following description and claims, the term “coupled” along with its derivatives, may be used. “Coupled” is used to indicate that two or more elements co-operate or interact with each other, but they may or may not have intervening physical or electrical components between them.

As used in the claims, unless otherwise specified the use of the ordinal adjectives “first”, “second”, “third”, etc., to describe a common element, merely indicate that different instances of like elements are being referred to, and are not intended to imply that the elements so described must be in a given sequence, either temporally, spatially, in ranking, or in any other manner.

The following clauses and/or examples pertain to further embodiments or examples. Specifics in the examples may be used anywhere in one or more embodiments. The various features of the different embodiments or examples may be variously combined with some features included and others excluded to suit a variety of different applications. Examples may include subject matter such as a method, means for performing acts of the method, at least one machine-readable medium including instructions that, when performed by a machine cause the machine to perform acts of the method,

or of an apparatus or system for facilitating hybrid communication according to embodiments and examples described herein.

Some embodiments pertain to Example 1 that includes an apparatus to facilitate dynamic merging of atomic operations, the apparatus comprising: detection/comparison logic, as facilitated by or at least partially incorporated into a memory controller, to detect atomic messages and a plurality of slot addresses; merging/conversion logic, as facilitated by or at least partially incorporated into a memory controller, to compare one or more slot addresses of the plurality of slot addresses with other slot addresses of the plurality of slot addresses to seek one or more matched slot addresses, wherein the merging/conversion logic, as facilitated by or at least partially incorporated into a memory controller, is further to merge the one or more matched slot addresses into one or more merged groups; and generation/execution logic to generate one or more merged atomic operations based on and corresponding to the one or more merged groups.

Example 2 includes the subject matter of Example 1, wherein the atomic messages comprise one or more of single instruction, multiple data (SIMD) atomic messages issued by an execution unit, wherein the execution unit comprises a graphics shader execution unit.

Example 3 includes the subject matter of Examples 1-2, wherein the memory controller is coupled to one or more processors having at least one of a graphics processor and an application processor.

Example 4 includes the subject matter of Examples 1-3, wherein memory controller comprises one or more of shared local memory (SLM) controller and data port controller.

Example 5 includes the subject matter of Examples 1-4, wherein the one or more merged atomic operations are based on retaining unique slot addresses and skipping recurring slot addresses.

Example 6 includes the subject matter of Examples 1-5, wherein the one or more merged atomic operations are forwarded on to a cache controller with atomic arithmetic logic unit (ALU).

Example 7 includes the subject matter of Examples 1-6, wherein the graphics processor is co-located with the application processor on a common semiconductor package.

Some embodiments pertain to Example 8 that includes a method for facilitating dynamic merging of atomic operations, the method comprising: detecting atomic messages and a plurality of slot addresses; comparing one or more slot addresses of the plurality of slot addresses with other slot addresses of the plurality of slot addresses to seek one or more matched slot addresses, wherein the one or more matched slot addresses are merged into one or more merged groups; and generating one or more merged atomic operations based on and corresponding to the one or more merged groups.

Example 9 includes the subject matter of Example 8, wherein the atomic messages comprise one or more of single instruction, multiple data (SIMD) atomic messages issued by an execution unit, wherein the execution unit comprises a graphics shader execution unit.

Example 10 includes the subject matter of Examples 8-9, wherein one or more of detecting, comparing, generating are performed by one or more components hosted by a memory controller of a computing device, wherein the memory controller is coupled to one or more processors having at least one of a graphics processor and an application processor.

Example 11 includes the subject matter of Examples 8-10, wherein memory controller comprises one or more of shared local memory (SLM) controller and data port controller.

Example 12 includes the subject matter of Examples 8-11, wherein the one or more merged atomic operations are based on retaining unique slot addresses and skipping recurring slot addresses.

Example 13 includes the subject matter of Examples 8-12, wherein the one or more merged atomic operations are forwarded on to a cache controller with atomic arithmetic logic unit (ALU).

Example 14 includes the subject matter of Examples 8-13, wherein the graphics processor is co-located with the application processor on a common semiconductor package.

Some embodiments pertain to Example 15 that includes a graphics processing system comprising memory coupled to a processor, the processor to: detect atomic messages and a plurality of slot addresses; compare one or more slot addresses of the plurality of slot addresses with other slot addresses of the plurality of slot addresses to seek one or more matched slot addresses, wherein the one or more matched slot addresses are merged into one or more merged groups; and generate one or more merged atomic operations based on and corresponding to the one or more merged groups.

Example 16 includes the subject matter of Example 15, wherein the atomic messages comprise one or more of single instruction, multiple data (SIMD) atomic messages issued by an execution unit, wherein the execution unit comprises a graphics shader execution unit.

Example 17 includes the subject matter of Examples 15-16, wherein one or more of detecting, comparing, generating are performed by one or more components hosted by a memory controller of a computing device, wherein the memory controller is coupled to one or more processors having at least one of a graphics processor and an application processor.

Example 18 includes the subject matter of Examples 15-17, wherein memory controller comprises one or more of shared local memory (SLM) controller and data port controller.

Example 19 includes the subject matter of Examples 15-18, wherein the one or more merged atomic operations are based on retaining unique slot addresses and skipping recurring slot addresses.

Example 20 includes the subject matter of Examples 15-19, wherein the one or more merged atomic operations are forwarded on to a cache controller with atomic arithmetic logic unit (ALU).

Example 21 includes the subject matter of Examples 15-20, wherein the graphics processor is co-located with the application processor on a common semiconductor package.

Example 29 includes at least one non-transitory or tangible machine-readable medium comprising a plurality of instructions, when executed on a computing device, to implement or perform a method as claimed in any of claims or examples 8-14.

Example 30 includes at least one machine-readable medium comprising a plurality of instructions, when executed on a computing device, to implement or perform a method as claimed in any of claims or examples 8-14.

Example 31 includes a system comprising a mechanism to implement or perform a method as claimed in any of claims or examples 8-14.

Example 32 includes an apparatus comprising means for performing a method as claimed in any of claims or examples 8-14.

Example 33 includes a computing device arranged to implement or perform a method as claimed in any of claims or examples 8-14.

Example 34 includes a communications device arranged to implement or perform a method as claimed in any of claims or examples 8-14.

Example 35 includes at least one machine-readable medium comprising a plurality of instructions, when executed on a computing device, to implement or perform a method or realize an apparatus as claimed in any preceding claims.

Example 36 includes at least one non-transitory or tangible machine-readable medium comprising a plurality of instructions, when executed on a computing device, to implement or perform a method or realize an apparatus as claimed in any preceding claims.

Example 37 includes a system comprising a mechanism to implement or perform a method or realize an apparatus as claimed in any preceding claims.

Example 38 includes an apparatus comprising means to perform a method as claimed in any preceding claims.

Example 39 includes a computing device arranged to implement or perform a method or realize an apparatus as claimed in any preceding claims.

Example 40 includes a communications device arranged to implement or perform a method or realize an apparatus as claimed in any preceding claims.

The drawings and the forgoing description give examples of embodiments. Those skilled in the art will appreciate that one or more of the described elements may well be combined into a single functional element. Alternatively, certain elements may be split into multiple functional elements. Elements from one embodiment may be added to another embodiment. For example, orders of processes described herein may be changed and are not limited to the manner described herein. Moreover, the actions of any flow diagram need not be implemented in the order shown; nor do all of the acts necessarily need to be performed. Also, those acts that are not dependent on other acts may be performed in parallel with the other acts. The scope of embodiments is by no means limited by these specific examples. Numerous variations, whether explicitly given in the specification or not, such as differences in structure, dimension, and use of material, are possible. The scope of embodiments is at least as broad as given by the following claims.

What is claimed is:

1. An apparatus comprising:

one or more processors including a graphics processor; and

a memory controller coupled to the one or more processors, wherein the memory controller is to:

detect atomic messages and a plurality of slot addresses of the atomic messages, wherein the plurality of slot addresses map to memory addresses;

generate a signal array to indicate unique memory addresses among the plurality of slot address;

detect, in the signal array, a number of instances of each of the unique memory addresses among the plurality of slot addresses;

generate an atomic operation corresponding to the atomic messages for each of the unique memory addresses; and

select one or more slot addresses of the plurality of slot addresses to drop based on the number of instances for each of the unique memory addresses.

51

2. The apparatus of claim 1, wherein the atomic messages comprise one or more of single instruction, multiple data (SIMD) atomic messages issued by an execution unit, wherein the execution unit comprises a graphics shader execution unit.

3. The apparatus of claim 1, wherein the one or more processors comprise an application processor.

4. The apparatus of claim 3, wherein memory controller comprises one or more of shared local memory (SLM) controller or data port controller.

5. The apparatus of claim 1, wherein the memory controller to generate the atomic operation comprises the memory controller to generate a merged atomic operation based on retaining unique slot addresses and skipping recurring slot addresses.

6. The apparatus of claim 1, wherein the atomic operation is forwarded on to a cache controller having an atomic arithmetic logic unit (ALU).

7. The apparatus of claim 3, wherein the graphics processor is co-located with the application processor on a common semiconductor package.

8. A method comprising:

detecting atomic messages and a plurality of slot addresses of the atomic messages, wherein the plurality of slot addresses map to memory addresses;

generating a signal array to indicate unique memory addresses among the plurality of slot address;

detecting, in the signal array, a number of instances of each of the unique memory addresses among the plurality of slot addresses;

generating an atomic operation corresponding to the atomic messages for each of the unique memory addresses; and

selecting one or more slot addresses of the plurality of slot addresses to drop based on the number of instances for each of the unique memory addresses.

9. The method of claim 8, wherein the atomic messages comprise one or more of single instruction, multiple data (SIMD) atomic messages issued by an execution unit, wherein the execution unit comprises a graphics shader execution unit.

10. The method of claim 8, wherein one or more of detecting, generating, selecting are performed by one or more components hosted by a memory controller of a computing device, wherein the memory controller is coupled to one or more processors having at least one of a graphics processor and an application processor.

11. The method of claim 10, wherein the memory controller comprises one or more of shared local memory (SLM) controller or data port controller.

12. The method of claim 8, wherein generating the atomic operation further comprises generating a merged atomic operation based on retaining unique slot addresses and skipping recurring slot addresses.

13. The method of claim 8, wherein the atomic operation is forwarded on to a cache controller having an atomic arithmetic logic unit (ALU).

14. The method of claim 10, wherein the graphics processor is co-located with the application processor on a common semiconductor package.

15. At least one non-transitory machine-readable medium comprising instructions that when executed by a computing device, cause the computing device to perform operations comprising:

detecting atomic messages and a plurality of slot addresses of the atomic messages, wherein the plurality of slot addresses map to memory addresses;

52

generating a signal array to indicate unique memory addresses among the plurality of slot address;

detecting, in the signal array, a number of instances of each of the unique memory addresses among the plurality of slot addresses;

generating an atomic operation corresponding to the atomic messages for each of the unique memory addresses; and

selecting one or more slot addresses of the plurality of slot addresses to drop based on the number of instances for each of the unique memory addresses.

16. The non-transitory machine-readable medium of claim 15, wherein the atomic messages comprise one or more of single instruction, multiple data (SIMD) atomic messages issued by an execution unit, wherein the execution unit comprises a graphics shader execution unit.

17. The non-transitory machine-readable medium of claim 15, wherein one or more of detecting, generating, selecting are performed by one or more components hosted by a memory controller of a computing device, wherein the memory controller is coupled to one or more processors having at least one of a graphics processor and an application processor, wherein the graphics processor is co-located with the application processor on a common semiconductor package.

18. The non-transitory machine-readable medium of claim 17, wherein the memory controller comprises one or more of shared local memory (SLM) controller or data port controller.

19. The non-transitory machine-readable medium of claim 15, wherein generating the atomic operation further comprises generating a merged atomic operation based on retaining unique slot addresses and skipping recurring slot addresses.

20. The non-transitory machine-readable medium of claim 15, wherein the atomic operation is forwarded on to a cache controller having an atomic arithmetic logic unit (ALU), and wherein the cache controller is coupled to one or more processors comprising a graphics processor and an application processor.

21. A system comprising:

a memory;

one or more processors including a graphics processor; and

a memory controller coupled to the memory and the one or more processors, wherein the memory controller is to:

detect atomic messages and a plurality of slot addresses of the atomic messages, wherein the plurality of slot addresses map to memory addresses of the memory;

generate a signal array to indicate unique memory addresses among the plurality of slot address;

detect, in the signal array, a number of instances of each of the unique memory addresses among the plurality of slot addresses;

generate an atomic operation corresponding to the atomic messages for each of the unique memory addresses; and

select one or more slot addresses of the plurality of slot addresses to drop based on the number of instances for each of the unique memory addresses.

22. The system of claim 21, wherein the atomic messages comprise one or more of single instruction, multiple data (SIMD) atomic messages issued by an execution unit, wherein the execution unit comprises a graphics shader execution unit.

23. The system of claim 21, wherein the one or more processors comprise an application processor.

24. The system of claim 23, wherein memory controller comprises one or more of shared local memory (SLM) controller or data port controller.

5

25. The system of claim 21, wherein the memory controller to generate the atomic operation comprises the memory controller to generate a merged atomic operation based on retaining unique slot addresses and skipping recurring slot addresses.

10

* * * * *